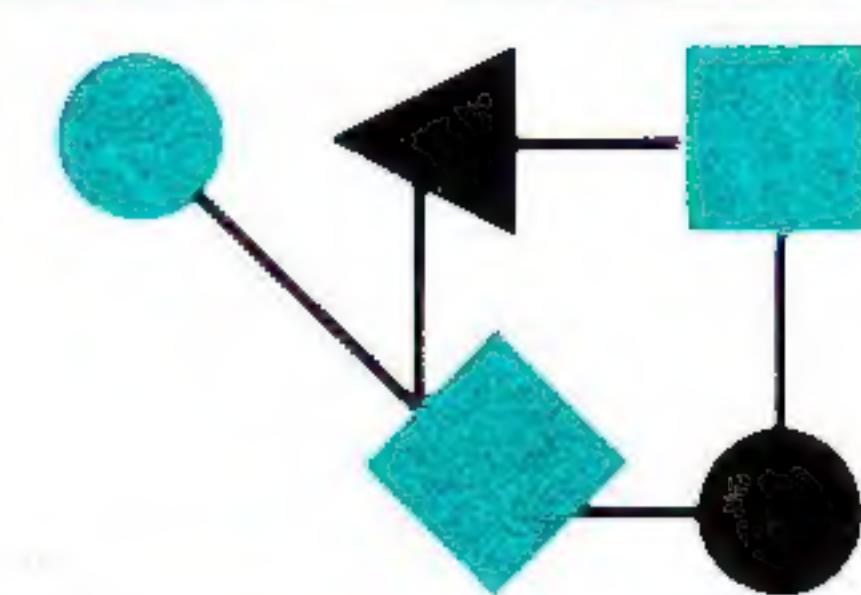


CONNE~~X~~IONS



The Interoperability Report

February 1990

Volume 4, No. 2

*ConneXions—
The Interoperability Report
tracks current and emerging
standards and technologies
within the computer and
communications industry.*

In this issue:

Interview with Steve Kent on Internet Security.....	2
The Trusted Mail System.....	8
COS/NIST testing program...	9
The INTEROP 89 network....	10
OSI at INTEROP 89.....	18
Letters to the Editor.....	25
Book Review.....	26

ConneXions is published monthly by Advanced Computing Environments, 480 San Antonio Road, Suite 100, Mountain View, California 94040, USA. Phone: 415-941-3399. Fax: 415-949-1779.

© 1990
Advanced Computing Environments.
Quotation with attribution encouraged.

ConneXions—The Interoperability Report and the ConneXions masthead are trademarks of Advanced Computing Environments.

ISSN 0894-5926

From the Editor

In the wake of the infamous *Internet Worm* of November 1988, a great deal of attention has been focused on network and system security. In the Internet community, much of the research and development in this arena is being done under the auspices of the *Privacy and Security Research Group* which is part of the Internet Research Task Force. This month we bring you an interview with the chairman of this group, Dr. Stephen Kent, Chief Scientist at BBN Communications. The interviewer is Daniel Dern.

The rest of this issue is mostly devoted to a look back at the INTEROP™ 89 show network. We do this with two articles, one describing the network design, and one reporting on the OSI demos which were featured at INTEROP 89.

The 1989 network was just over double the size of the one featured at the 1988 show. While show management took responsibility for the main network design, the actual setting up, trouble-shooting and maintenance of the show network was performed with the invaluable aid of a team of networking experts from many different companies and research institutions in the industry. Steve Knowles, the author of our first article, was one of the lead engineers on that force of volunteers.

For those of you who keep statistics, the main show network backbone consisted of almost six miles of cabling (not counting the local cabling inside each exhibitor's booth), and a two-block link of 10Mbit Ethernet over line-of-sight microwave to the Fairmont Hotel, (in turn connected to UTP Ethernet on the remote side). It took a grand total of around 1400 man-hours of labor on the part of the shownet team to create the whole multi-vendor, show-wide network. The resulting network had about 600 end-systems (in 101 booths) connected to it, and there were at least 10 connections to the outside world from the floor.

Dave Katz from Merit/NSFNET spent a great deal of time during INTEROP 89 working on the OSI demonstrations. In an article starting on page 18 he gives an overview of these demos and talks about the important lessons learned.

Also in this issue, you will find some news about the recently announced COS/NIST agreement to operate an OSI testing program, some letters to the Editor, and a review of an ISDN book.

The 1989 ConneXions Table of Contents sheet is now available. It was mailed out to all subscribers with the January 1990 issue. Similar sheets for 1987 and 1988 are also available free of charge.

Interview with Steve Kent on Internet Security

by Daniel P. Dern

Introduction

Security for our computers, networks, and the data that inhabits them, has gotten increasing attention in the past year or two, what with the Chaos Club crackers nabbed by Cliff Stoll, the Internet and WANK worms, and other incidents. For the Internet community, our ability to use it is only as good as our ability to trust it. For a look at what's in the works to increase the security levels of Internet activity, we turned to Dr. Stephen Kent, Chief Scientist at BBN Communications, and head of the Internet Research Task Force's Privacy and Security Research Group (PSRG).

What does security mean, in terms of the Internet, currently?

I think it's useful to divide security issues, in the Internet environment especially, into three categories:

- Issues associated with protection of subscriber traffic;
- Issues related to protection of transmission and switching resources;
- and issues associated with protecting access to end user resources, i.e., hosts attached to the network.

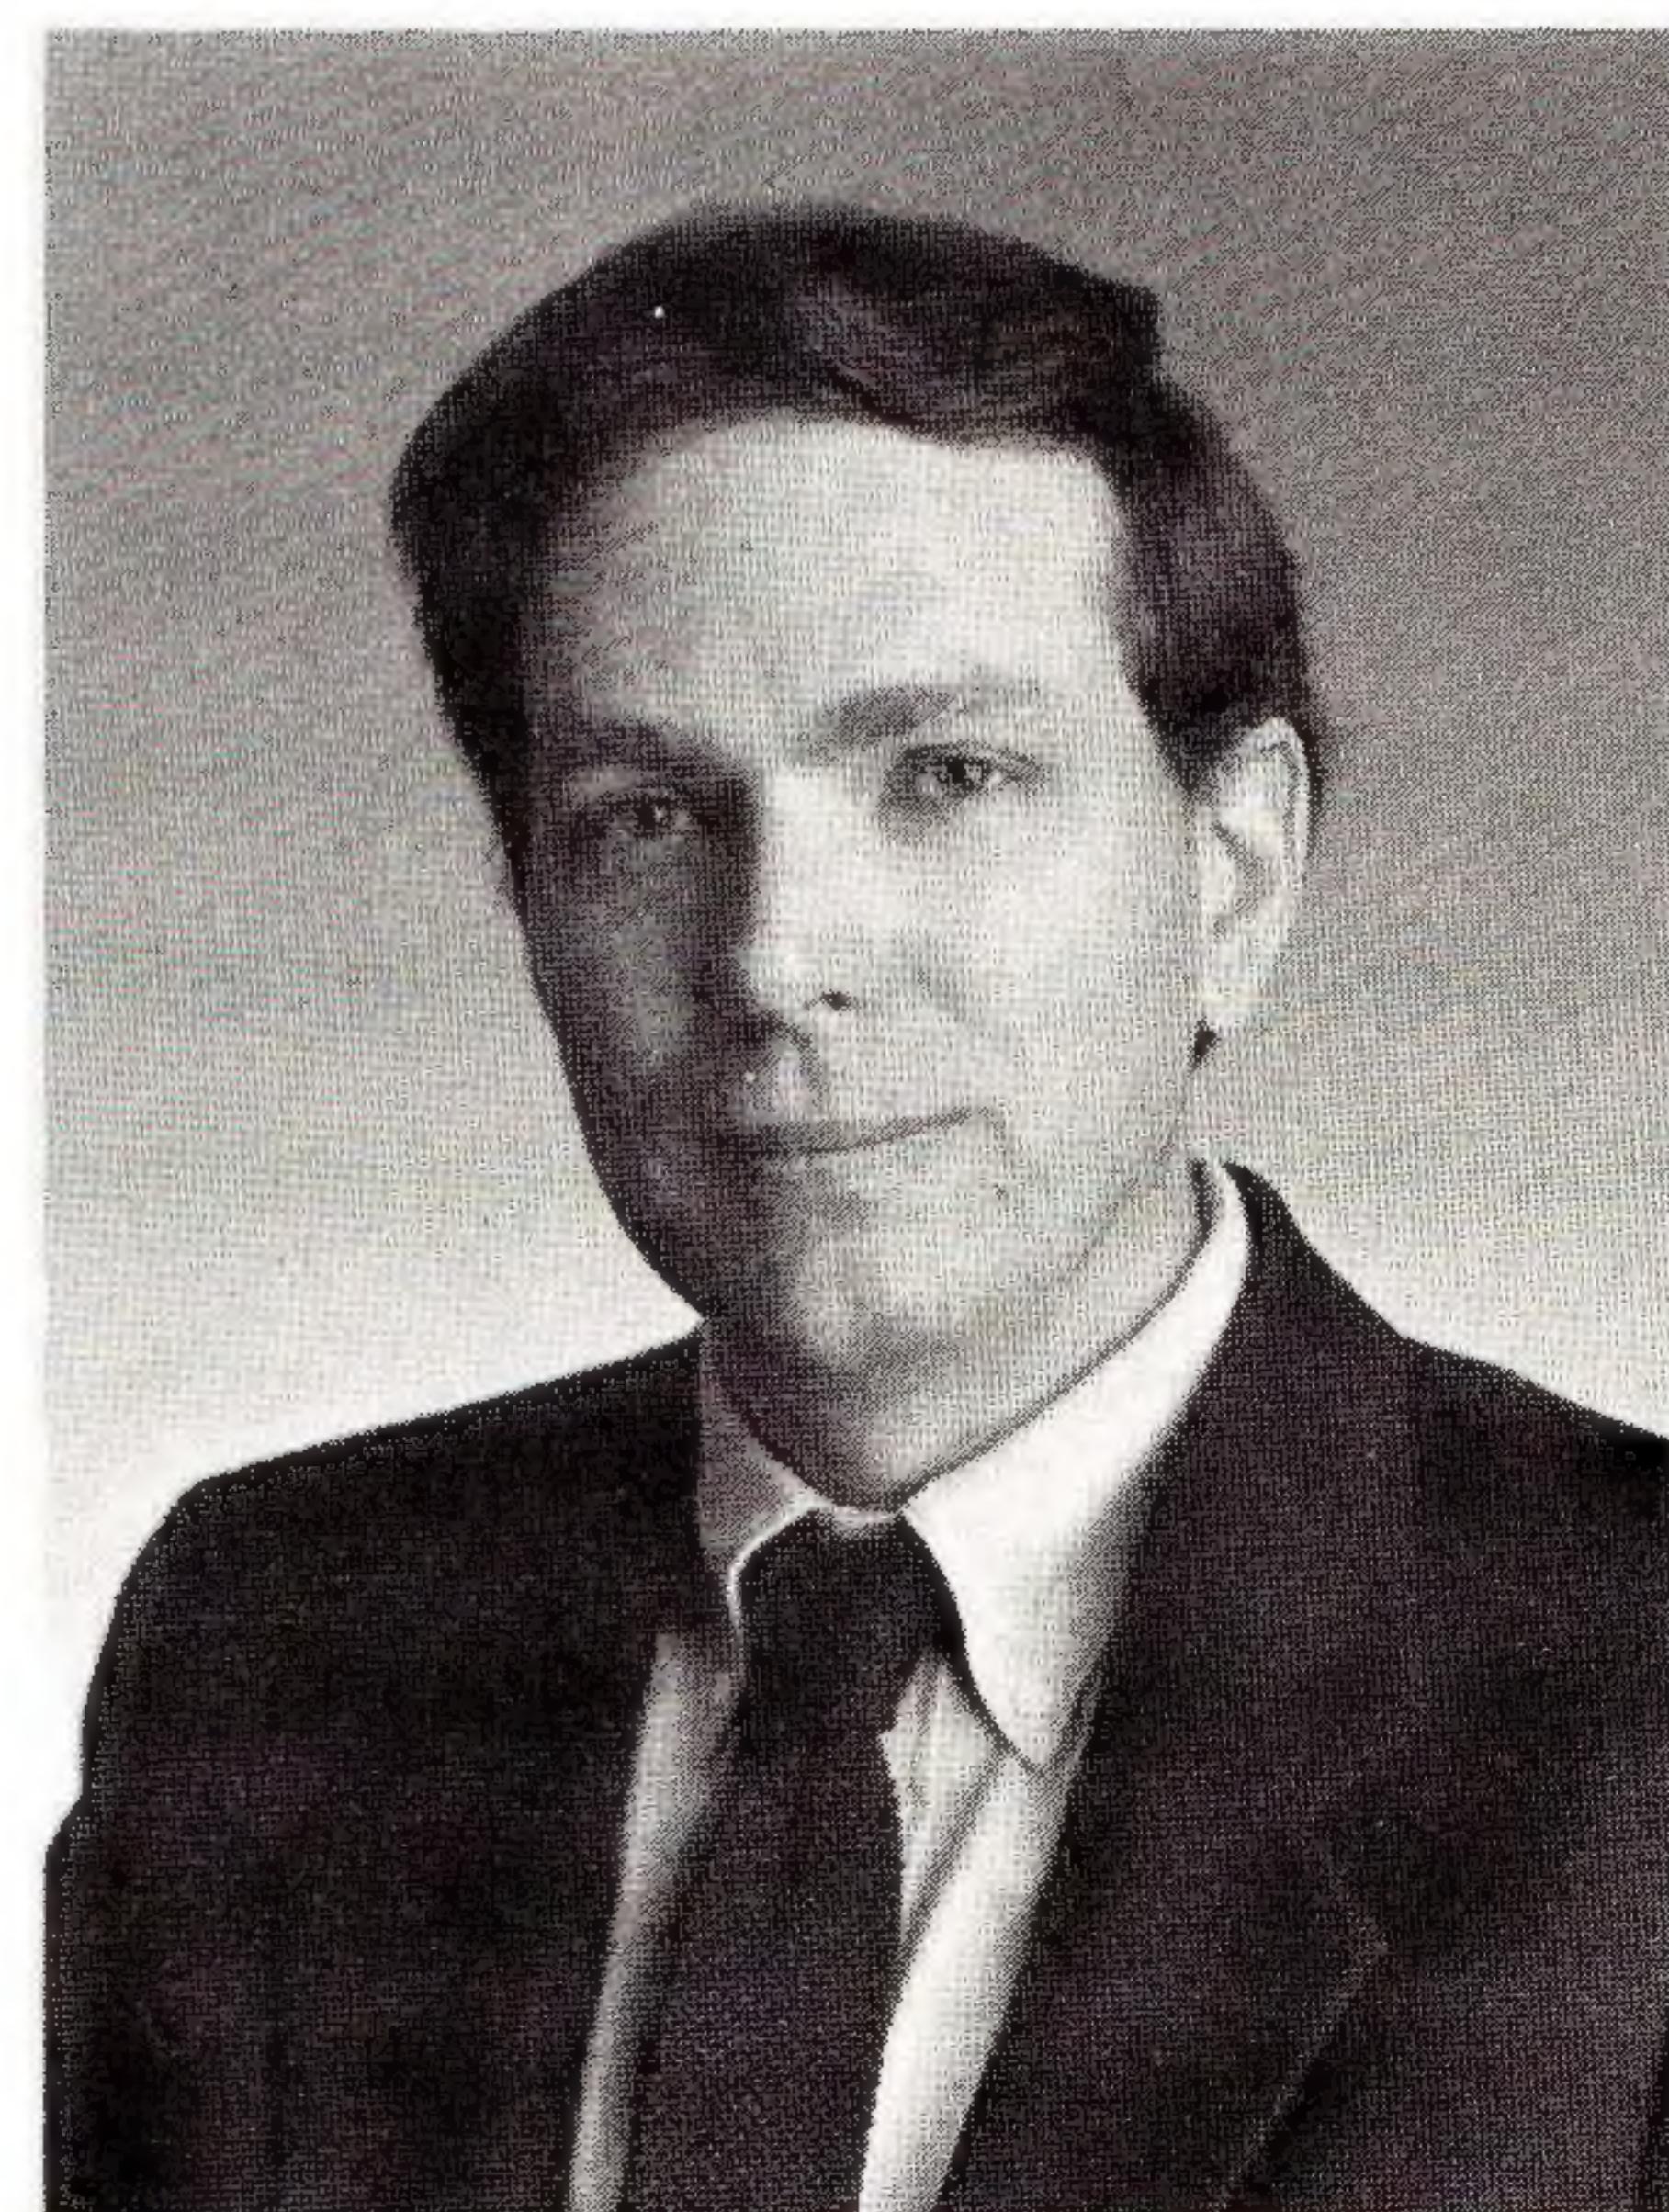
It turns out that most people now focus on the third category. People are very concerned about the increased vulnerability of their computers as the result of being hooked up to any network, not just the Internet. This is a justifiable concern.

Administration

But in most cases the concern really should be focused on the lack of good administrative practices and security software in subscriber computers, rather than with any intrinsic security shortcomings of the Internet.

Consider, for example, the *Internet Worm* incident from a year ago. In this incident the unauthorized access to subscriber computers was not caused by any security flaws in the Internet switching and transmission resources per se, nor in standard Internet protocols as specified in RFCs.

What allowed the worm to propagate was a combination of sloppiness, essentially...first on the part of some vendors, and secondly on the part of host administrators, who didn't enforce good password management standards.



What are some of the solutions to security problems like those we're currently experiencing?

Good local security administration, combined with better quality host software—software that isn't distributed with debug switches on, for example,—would go a long way toward dealing with many of the security issues that people face.

Use of evaluated operating systems, ones which have been examined and rated by the National Computer Security Center, would also be a step in the right direction.

Network security measures can, in principle, extend host-based security measures through the Internet or can be used to bolster host defenses in various ways.

Such as?

Traffic filtering

Well, some sites already use gateways to filter inbound and outbound traffic, to restrict connectivity between local hosts and those elsewhere on the Internet. This can be a good first step in controlling access to the local hosts.

However, it also is important to remember that, in a network environment, security for an individual host often cannot be achieved only by having good software in that host and through good, local password management. Security for that host may often depend upon the security of the other hosts with which it communicates. Gateways are limited in their ability to provide fine-grained control of access to hosts due in part to the susceptibility of network addresses to forgery. End-to-end encryption can be used to address this problem, both for realtime and staged delivery communication applications.

What is being done in the Internet as relates to this?

Privacy Enhanced Mail

The IRTF's Privacy and Security Research Group (which I chair) has been working in a number of directions over the past several years. One of them has been a project called *Privacy Enhanced Mail*, addressing the problem of secure electronic mail. [As described in RFCs 1113, 1114 and 1115].

Why this particular focus?

We felt it was something that was needed, as a retrofit to the existing Internet protocol environment.

Have there been reports of users feeling like they have been "spoofed" by fraudulent mail, had their mail read otherwise intercepted? Or is it simply time to be more rigorous about network security?

There are a couple of reasons for security that apply particularly to e-mail across networks. Some people have observed that they send information by Internet mail that they really feel uncomfortable sending, and would feel much more comfortable if it were protected. Others say there are types of information they *would* send, if only the mail *were* protected. So we need more security both for current traffic, and to permit fuller use of network mail.

What are some of the significant aspects of the Privacy Enhanced Mail project?

The cryptography we make use of is designed to shift the burden so you can have a high degree of confidence in the mail you receive that it came from who it says it does, without having to have a high level of trust in all the gateways, switches, and mail relay hosts along the path between originator and recipient.

continued on next page

Interview with Steve Kent (*continued*)

How much of this might apply not only to the Internet but to the non-Internet community?

Ideas from X.400

The literal, detailed mechanisms used in Privacy Enhanced Mail are tied to the Internet protocol suite, but the concepts that underlie these mechanisms are more general. For example, 1988 X.400 provides an analogous and broader set of security facilities based on the same basic mechanisms. So if subscribers are using X.400, it doesn't make sense to ask if they can apply our work to that mail system. It is worth noting that the "secure" mail mechanisms for both the Internet and X.400 use the same type of public key certificates.

Once we have this system set up for Internet users, it will be easier for them to transition into (secure) X.400 as it becomes available, because the key management infrastructure will be already established. That's a big part of the overall system and it's essentially the same in both protocol suites. That was a conscious decision on our part, by the way, to not go in a different direction from existing standards where possible.

How does Privacy Enhanced Mail work?

Briefly, we begin by computing a checksum-like value called a *one-way hash* on the email message and we encrypt the result with the private RSA key of the sender. Then we generate a *Data Encryption Standard* (DES) key and use it to encrypt the message. We then use the public RSA key of the recipient to encrypt the DES key used to protect this message. The recipient reverses this process to decrypt the message and validate the identity of the originator.

RSA

RSA is a public-key cryptosystem developed in 1977 by MIT Professors Ronald Rivest, Adi Shamir, and Len Adleman. The RSA algorithm is itself too slow to apply to the text of messages, particularly when larger keys are used. However, it's fast enough to encrypt the DES key, one copy of which is prepended to the message for each recipient.

What is the role of the RSA technology?

Key management

What RSA gives us is a tool around which we developed a secure public key distribution scheme. This is the hardest part of security, in many ways—establishing a workable "hierarchy of trust" to enable a user to easily validate the identity of other users as originators or recipients of messages.

A major challenge for key management systems for public key cryptography is the reliable association of a user's public key with his identity. In general, you can't just look it up in an on-line directory system. You probably should not trust that system, and even if you could trust it, you would require a secure (tamper-proof) communications path between you and that directory system. Without strong security assurances about the directory system and the path to that system you could wind up with some attacker's public key instead of the public key for the person with whom you want to communicate. Providing a secure, practical means of associating a user's identity with his public key is absolutely critical to general use of public key technology.

Public Key Certificates

The solution to this problem is what we call *Public Key Certificates*. A certificate is a data structure which includes the ID of the “subject,” his public key, a *validity interval*, and the ID of the “issuer” of the certificate. The subject is the entity who knows the private key corresponding to the public key contained in this certificate. The validity interval is like the pair of dates on many credit cards, specifying the time interval during which the certificate is valid. The issuer field identifies the person or organization who vouches for the binding between the public key and the subject’s ID. This data structure is digitally signed by the issuer using his private key, thus providing the requisite binding.

In our privacy enhanced mail system, the originator passes his certificate in the message header and he may optionally include the certificate for the issuer of his certificate. When a recipient gets a message he uses the public key of that issuer to validate the originator’s certificate and thus get the originator’s public key. To validate the issuer’s certificate the recipient needs the public key of whoever issued that certificate, and so on recursively.

Hierarchy of certificates

This approach naturally generates a certificate hierarchy which, should map to a trust relationship. People speak about “walking down a chain of certificates” to validate a specific user’s certificate. We have constrained the hierarchy for the privacy enhanced email system to minimize the length of these chains so as to simplify the validation process. We also expect users to cache validated certificates to further reduce the time required to process these protected messages.

You need a trusted place to start from, and to traverse to.

Yes. It’s very important that the relationship between the subject and the issuer of a certificate be such that a user will believe that the issuer should be trusted to vouch for the identity of the subject. In our scheme we have constrained the certification hierarchy to match the naming hierarchy as employed in X.500. This is more restrictive than the X.509 specification, but we feel it is an appropriate constraint.

You could use the same key management technology to distribute keys and perform authentication for other applications as well—real time end-to-end encryption for TCP connections, for example. It could also be used for authenticating gateway-to-gateway communication.

Remember that the public key certificate infrastructure adopted for privacy enhanced mail is taken right out of CCITT X.509. The PSRG did not invent the concepts here, we only adopted some constraints and specified some administrative and procedural aspects that we felt were appropriate. That’s why I expect we’ll see these certificates used increasingly in networking communities, because the authentication framework really is an international standard, and is thus attractive in that regard.

How does this compare to other systems, say, Kerberos?

Kerberos

Kerberos has an advantage from a very pragmatic standpoint—MIT is ready to provide it to users, the same way they provide X Windows.

Interview with Steve Kent (*continued*)

But the Kerberos system is actually relatively old technology from a cryptographic key management standpoint. It's a technology that was described in the open literature and implemented in experimental systems over a decade ago. This is not a criticism of Kerberos per se, but public key certificates represent a more modern technology and one which is becoming very popular outside the U.S.

Centralized management

Also, the more centralized management aspects of Kerberos require that a user trust machines at various locations on a network to keep secret his key. Kerberos requires multiple (trusted) servers for a robust implementation, as opposed to the distributed, but not trusted, directory servers used in a certificate-based system. Different administrative domains tend to require separate servers because of these trust issues. Through bi-lateral agreements between domains Kerberos can provide interoperability across these boundaries, but, again, servers in each domain must be involved, in real time, to support inter-domain authentication.

It's more complicated to manage?

I have the sense is that it can become fairly complex. The X.509 approach has some fairly natural delegation procedures embodied in it, in terms of who trusts whom, that should make it easier to manage and more robust. But the Kerberos folks have real, practical experience with their system, and we need to see if the advantages that public key certificates promise turn out to be real in practice. There is one clear advantage of the certificate scheme; only the individual user knows his private key. Nobody else ever has to know. As a user, I'd feel much more comfortable about that.

Where else will we see security being added to our networks?

Filtering on IP header fields

Increasingly, we'll see gateways used for traffic filtering based on IP header fields and TCP or UDP port fields. By watching these header fields in packets, you can have a gateway permit/deny access based on criteria such as source and destination addresses, directionality, and even the nature of the application—e.g., remote login versus mail.

Gateways are a good place to manage internetwork traffic. They tend to be resources that are controlled by a (local or campus) network administrator and thus more amenable to implementing security services than all of the workstations and hosts on a LAN. It's a very attractive thing to do, for a variety of technical and administrative reasons.

Given what we're talking about in terms of security, and its various aspects, are there any quantifications of what it costs, per user, site, network, etc.?

Cost

Cost has always been the big bugaboo in security, particularly in the unclassified environment. The things we're talking about do have some definable costs. Gateway filtering, if it becomes part of the gateway requirements RFC, should become a "built in" feature that all gateways exhibit, so there won't be any additional capital cost for this functionality. I think we are reaching a point where it won't cost you any more to buy a gateway that has this filtering in it.

There may be some performance "hit," but this can be kept small, if you're clever. And there will be some administrative overhead, relating mostly to how complex your network environment is, how often people want to change security tables, etc.

Reference implementation

As for Privacy Enhanced Mail, this capability can be added on a software basis to Internet hosts, and we will be distributing free copies of the software to the Internet community in the Spring of 1990. Of course vendors may offer feature-laden, supported versions for sale in addition to the free distribution. But there will be a reference implementation distributed in C, to work under UNIX, available to anybody in the Internet. The cost then becomes per mailbox, if you will, on a two year basis, arising out of the licensing agreement with RSA Data Security Inc. for the use of the RSA algorithm, plus their involvement in the key management infrastructure.

What kind of performance toll will the Privacy Enhanced Mail software add to user and system activities?

It's not clear. But we hope it will be reasonably painless. Remember that we're only doing this initially for mail, which doesn't require real-time interactive response. The encrypting can be done as a background task once the message has been composed and submitted for privacy processing. The processing overhead is a function of the message length and the number of recipients.

In software, we can do encryption at rates of 50–60 kilobits per second and higher, depending on the speed of your workstation. The RSA operations take a few seconds per recipient. That's a negligible delay—and well worth it, if it lets you use e-mail instead of more time-consuming alternatives.

References

Linn, J., "Privacy enhancement for Internet electronic mail: Part I—message encipherment and authentication procedures." [Draft], RFC 1113 (Obsoletes RFC 989, RFC 1040).

Kent, S.T.; Linn, J., "Privacy enhancement for Internet electronic mail: Part II—certificate-based key management." [Draft], RFC 1114.

Linn, J., "Privacy enhancement for Internet electronic mail: Part III—algorithms, modes, and identifiers." [Draft], RFC 1115

Schiller, J., "Kerberos: Network Authentication for Today's Open Networks," *ConneXions*, Volume 4, No. 1, January 1990.

CCITT Recommendation X.509/ISO Standard ISO-9594-8 "The Directory—Authentication Framework."

[Ed.: See also "The Trusted Mail System" on the following page].

DANIEL P. DERN is a Massachusetts-based independent writer specializing in technology and business. He also writes computer humor, musical theater, science fiction, and how-to's on consulting and PR. During his six years at BBN as PR Manager, marketing writer, and technical writer, he wrote about ARPANET and UNIX technology, including the DDN's Non-Technical Network Overview, and numerous feature articles for the trade press.

The Trusted Mail System

by Daniel P. Dern

The initial implementation of Internet Privacy Enhanced Mail is being worked on by organizations including BBN Communications Corporation, Cambridge, Mass.; RSA Data Security Inc., Redwood City, CA.; and Trusted Information Systems (TIS), Glenwood, MD.

TMail

According to a recent posting to the TCP-IP mailing list, "TIS is implementing the Trusted Mail System (TMail), as part of a DARPA sponsored research project. TMail is a prototype privacy enhanced mail system designed to investigate embedded cryptography within a trusted system. It was based on the message processing procedures described in RFC 1040.

At the request of DARPA and the Privacy and Security Research Group, TIS is planning to modify the TMail system to be compliant with the recently released RFC 1113, RFC 1114 and RFC 1115 and to make it available to the Internet. TIS has been working closely with RSADSI and BBN with respect to the release of TMail and other software necessary to support message processing and the key management procedures described in the RFCs."

TMail uses DES for message authentication and encryption, and RSA technology for key management and digital signature purposes. The TMail system has been incorporated into the Rand Message Handling System, and runs on Sun Workstations using SunOS (UNIX).

TIS expects to simplify the process of porting TMail to other mail systems, by separating the MH-specific aspects of the program from the non-mailer specific ones.

Getting TMail

TMail is currently undergoing Alpha testing at TIS, to be followed by Beta testing by PSRG members. "After testing in that community and after we have completely integrated it with the other software from RSADSI and BBNCC, and tested that, it will be made available on a wider basis to the general Internet," according to TIS Scientist and IETF PSRG member David Balenson. "Distribution will have to take into account issues like export control, since it will contain cryptographic software. So it will probably *not* be available through general anonymous FTP."

For further information on TMail, contact:

Pamela Cochrane
Trusted Information Systems
3060 Washington Road (Rt. 97)
Glenwood, MD. 21738
301-442-1673

cochrane@tis.com

COS and NIST establish Testing Program

The National Institute of Standards and Technology (NIST) and the Corporation for Open Systems International (COS) have signed an agreement to help prepare for the introduction of new computer networking standards in the federal government.

NIST is establishing a testing program to ensure that networking products purchased by federal agencies comply with the *Government Open Systems Interconnection Profile*—GOSIP (Federal Information Processing Standard 146). After August 15, 1990, federal agencies must use the GOSIP specifications in procuring networking products.

GOSIP

GOSIP defines a set of data communication rules called "protocols" which enables computer systems developed by different vendors to communicate, and enables the users of different applications on these systems to exchange information. The backbone of GOSIP is the internationally accepted *Open Systems Interconnection* (OSI) protocols. The technical specifications for GOSIP were developed in concert with vendors and users of computer networks in workshops organized by NIST.

Products based on GOSIP are being developed but methods are needed to test them for conformance to the GOSIP specifications and to test whether different manufacturers' products in fact can communicate with each other.

COS Test Suite

COS is a not-for-profit consortium formed in 1986 to accelerate the introduction of interoperable, multivendor products and services based on OSI, ISDN (*Integrated Services Digital Networks*) and related standards. COS has developed a number of test suites and testing tools to help vendors design and test systems that meet standard OSI specifications for interoperability.

As part of the agreement, COS will provide technical support and several test suites for conformance testing which will be evaluated by NIST.

"We welcome the opportunity to work with the experienced staff at COS. Although much remains to be accomplished, we feel confident that by the time the August deadline rolls around, policies, and procedures will be in place for testing compliance with GOSIP," says Kevin Mills, chief of the NIST Systems and Network Architecture Division.

"The GOSIP protocols are an important step in the widespread introduction of OSI products. This agreement gives COS an opportunity to contribute to the conformance testing infrastructure so necessary to a successful GOSIP compliance program," says Lincoln D. Faurer, president and chief executive officer of COS.

NVLAP

Also as part of the program, NIST will develop criteria by which its *National Voluntary Laboratory Accreditation Program* (NVLAP) will accredit outside laboratories to evaluate test systems as well as other laboratories which test vendors' implementations of GOSIP.

The joint COS/NIST venture is expected to last through September 1990.



Introduction

The INTEROP 89 Network: From one of its builders

by Stev Knowles, FTP Software

There are infinite reasons to go to the INTEROP trade show. Daniel Dern went through several in his *ConneXions* article in the November 1989 issue. In the past, when I have gone to INTEROP, I have done a wide variety of things, going to class sessions to learn about new protocols my company had an interest in, attended BOF sessions concerning issues from the *extremely* boring to the *excitingly* fascinating. INTEROP 89, so far the biggest INTEROP ACE has had, held the promise of being a place to meet the people I rarely see, and to discuss (allright, argue) about current hot topics. Unfortunately, I didn't get much of a chance to do *any* of these things. Fortunately, what I *did* get to do was infinitely more interesting than arguing about routing with people who already have their minds made up.

Before I get lost in detail, allow me to clarify some points. I am *sure* that there will be people whom I will not mention here, who deserve some credit for their assistance in INTEROP 89. The number of companies and individuals who provided hardware and assistance is monumental. Rarely have I seen so many companies help each other out, with loans of equipment and technical expertise to bring everyone's networks up with a minimum amount of hassle and stress. Several companies provided equipment from their booths for the INTEROP show to use for emergency shortages and changes in topology, some just hours before the doors opened Wednesday at noon. The network ran smoothly, once the show opened. The only complaint that made its way back to us was one conference attendee who felt that everything was running *too* smoothly. There was no-one running around, typing hurriedly at a terminal, while 5 people stood around him with concerned expressions on their faces. Whoever you were: sorry about that!

The Real Heroes

There are two people who deserve most of the credit for the success of the INTEROP show network, Peter de Vries, and Lisa Robertson, both from ACE. These are the people who spent months planning out the physical topology, adding up the lengths of cable runs, locating fan-out boxes and concentrators, coordinating with the vendors who loaned all the hardware that makes a network run, making sure that everyone was listed in all the appropriate lists for their network drops, and making sure that the 25 or 30 cases of assorted soda that the technical people consumed were always cold. This is an important point: if they remembered to keep the soda cold, you can assume that they had thought of everything else too. They had.

You may wonder why I am writing this, if all the credit for this goes to Peter and Lisa. Well, to be honest, so do I. I was invited to join about 10 other people to put together, play with, and take apart an arbitrarily complex network. A lot of people take this very seriously, and ignore that this is in reality a challenging goal. During the brief time it is up, it is a pretty amazing toy. The fact it was kept running in a near flawless condition was due to our pride in its magnificence. Our grand toy should appear nothing less than awe-inspiring to other people. As for our motivation for putting forth all this effort, if you have to ask, you wouldn't understand the answer.

This article is not intended to give one a coherent story about what went on in creating the shonnet for the INTEROP 89 show, but rather to provide highlights of the difficulties we ran into, and explanations of the options as we saw them, and our decisions and reasoning behind the decisions. It is possible (actually, probable) that any given decision could have been made better, but one must realize that the people who produced the shonnet did so in 4 days, working 18 hour days.

An early start

Friday morning started early; 5:30am wake up call, 6:00am meeting. Valarie Collins (ACE) had reserved a room for us every morning at 6:00am for breakfast until Wednesday. The first breakfast included about 20 people, including reps from some of the vendors who were loaning us hardware and some of the more esoteric other toys we needed (fiber optic cable, breakout boxes, line testing equipment).

Within 3 hours it became clear that there were about 10 serious people here. They were, in no particular order, Peter, Lisa, Steve Larbig (ACE), Dave Bridgham (Epilogue Technology), Karl Auerbach (Epilogue), John Romkey (Epilogue), Eric Brunner (Tule Network Services), Alan Brunner (consultant), Shelly de Vries (ACE), and myself (FTP Software). These are the people who *understood* why we were there, and are the people who put in 18 hour days from Friday until Wednesday afternoon. Some of the other people who should be mentioned are Geoff Baehr (Sun), Jeff Burgan (NASA-Ames), Ronnie Hueter, Phil Almquist (Stanford University) and Chris Lynch (ACE).

Network drops

Friday morning was spent marking the approximate locations of the network drops in the booths (some kind souls had already marked the booth locations for us). We just did a rough estimate of the locations, since they would be hanging from the ceiling some 30 feet above us, we could afford a bit of leeway in the exact locations that the wires had to be dropped to. Peter had spent some serious time involved in network planning, when we arrived there were maps of everything we could possibly want. There were maps for all the major media types, UTP (Unshielded Twisted Pair), Thin Ethernet, Thick Ethernet, 802.5, and FDDI fiber. These maps were further sub-divided into "work sheets." Work sheets were parts of the master maps in smaller, workgroup sized allotments. The long, "backbone" thick Ethernet runs were on one map, the thin Ethernets fanning out from each fan-out location were on others.

Media like UTP and thick Ethernet which had a large number of drops were put on maps based on sections of the show floor (east, middle, and west). Copies of these maps were used by crews of people who went off to do their assigned tasks. When they were done, and when time permitted, someone else would go over the work they had done to make sure everything was in place. As it turned out, we still ended up missing about 3 network drops. Considering there were several hundred, this, I suppose, should not reflect badly on us.

Our greatest resource was having some people on tap who have been doing this for a *long* time, and along with being up on all the latest technology, are pretty flexible people, willing to scrap the well laid foundation Peter gave us and strike out on our own when it appeared that our previous plans were not going to work. Fortunately, Peter was as agile as we were, and changes were quickly hashed out and implemented.

continued on next page

The INTEROP 89 Network (continued)

Running wires out from the exhibit hall

We found, as we were running wires, that we had been given incomplete information about the building we were working in. We had been told that one could get into the ceiling and thereby cross from the exhibit hall to the main concourse running along the front of the building. This we anticipated using to provide cable runs to the mail centers (one each on the east and west ends of the convention center), the microwave relay (for the mail center in the Fairmont Hotel, across the street), and a set of information terminals being provided by Prime for the concourse area. We found that the exhibit hall's cinder block wall went to the roof, and was sealed to the roof with mortar to provide fire walls. We sent off people to detail the building, and to provide recommendations for how to get the cable out of the exhibit hall and into the concourse area in light of this obstacle. We were brought back several suggestions:

- 1) We could just run the wires out one of the doors, and then run them back to the ceiling in the hall. The benefit to this approach was that it would be quick to implement. The problems were many, as we saw it. Since we were running cable through a door, it would have to be left open all day and night, requiring another security guard to sit by it. We decided that while we could put a connector in the cable at this point, and disconnect it at night, that would be cutting off email service after show hours. We also decided we could not use one of the main doors, since people would tend to pull on the wires and such. We also decided it was not aesthetically pleasing.
- 2) We could punch a *small* hole through the cinder block wall. This was eventually decided to be anti-social (!)
- 3) The convention center includes floor access plates for power and phone all over the show floor. These open into the garage under the convention center. The people we sent back out were to determine how to get the cables from the garage to the drop locations along the concourse in the front of the building. Unfortunately, there was no access between the garage and the concourse that didn't involve propping open a door. We decided it would be worse to leave these wires hanging out so far from us, and we didn't want to leave a guard down with them.
- 4) The final suggestion, and the eventual solution involved careful study of how a large building is put together, and playing on the weaknesses that show up. The convention center is not *really* one building, it is actually 4 buildings built 5 inches apart. The space between them is covered by aluminum panels. This allows the building to expand and contract as it heats and cools. As it turns out, you can also get from the exhibit hall to the concourse through these passages. The problem here is that these expansion joints tend to travel under walls.

Expansion joints

The convention center breaks the large hall up into 4 pieces, each slightly larger than the preceding one. This is why the concourse in front of the building is much larger on the east side than the west side, since the walls of the exhibit hall move farther towards the front of the building as you go the the west (see Figure 1). The building joints tend to appear at the locations where the walls jut out to make the next hall bigger.

We eventually managed to pass cables through the expansion joints, under the walls, and into the concourse area. This involved an incredible amount of contortion, but seemed the best solution to the problem. We could then continue the run under the floor (in the expansion joint), to the front of the building, where we planned to have the mail center runs, and the microwave and prime terminals were set up. We ended up spending a lot of time trying to figure this one out, and it involved some careful, realtime planning to allow us to remove expansion joint covers when the convention center was not being used by INTEROP or other tenants (this usually happened after midnight or 1 am.).

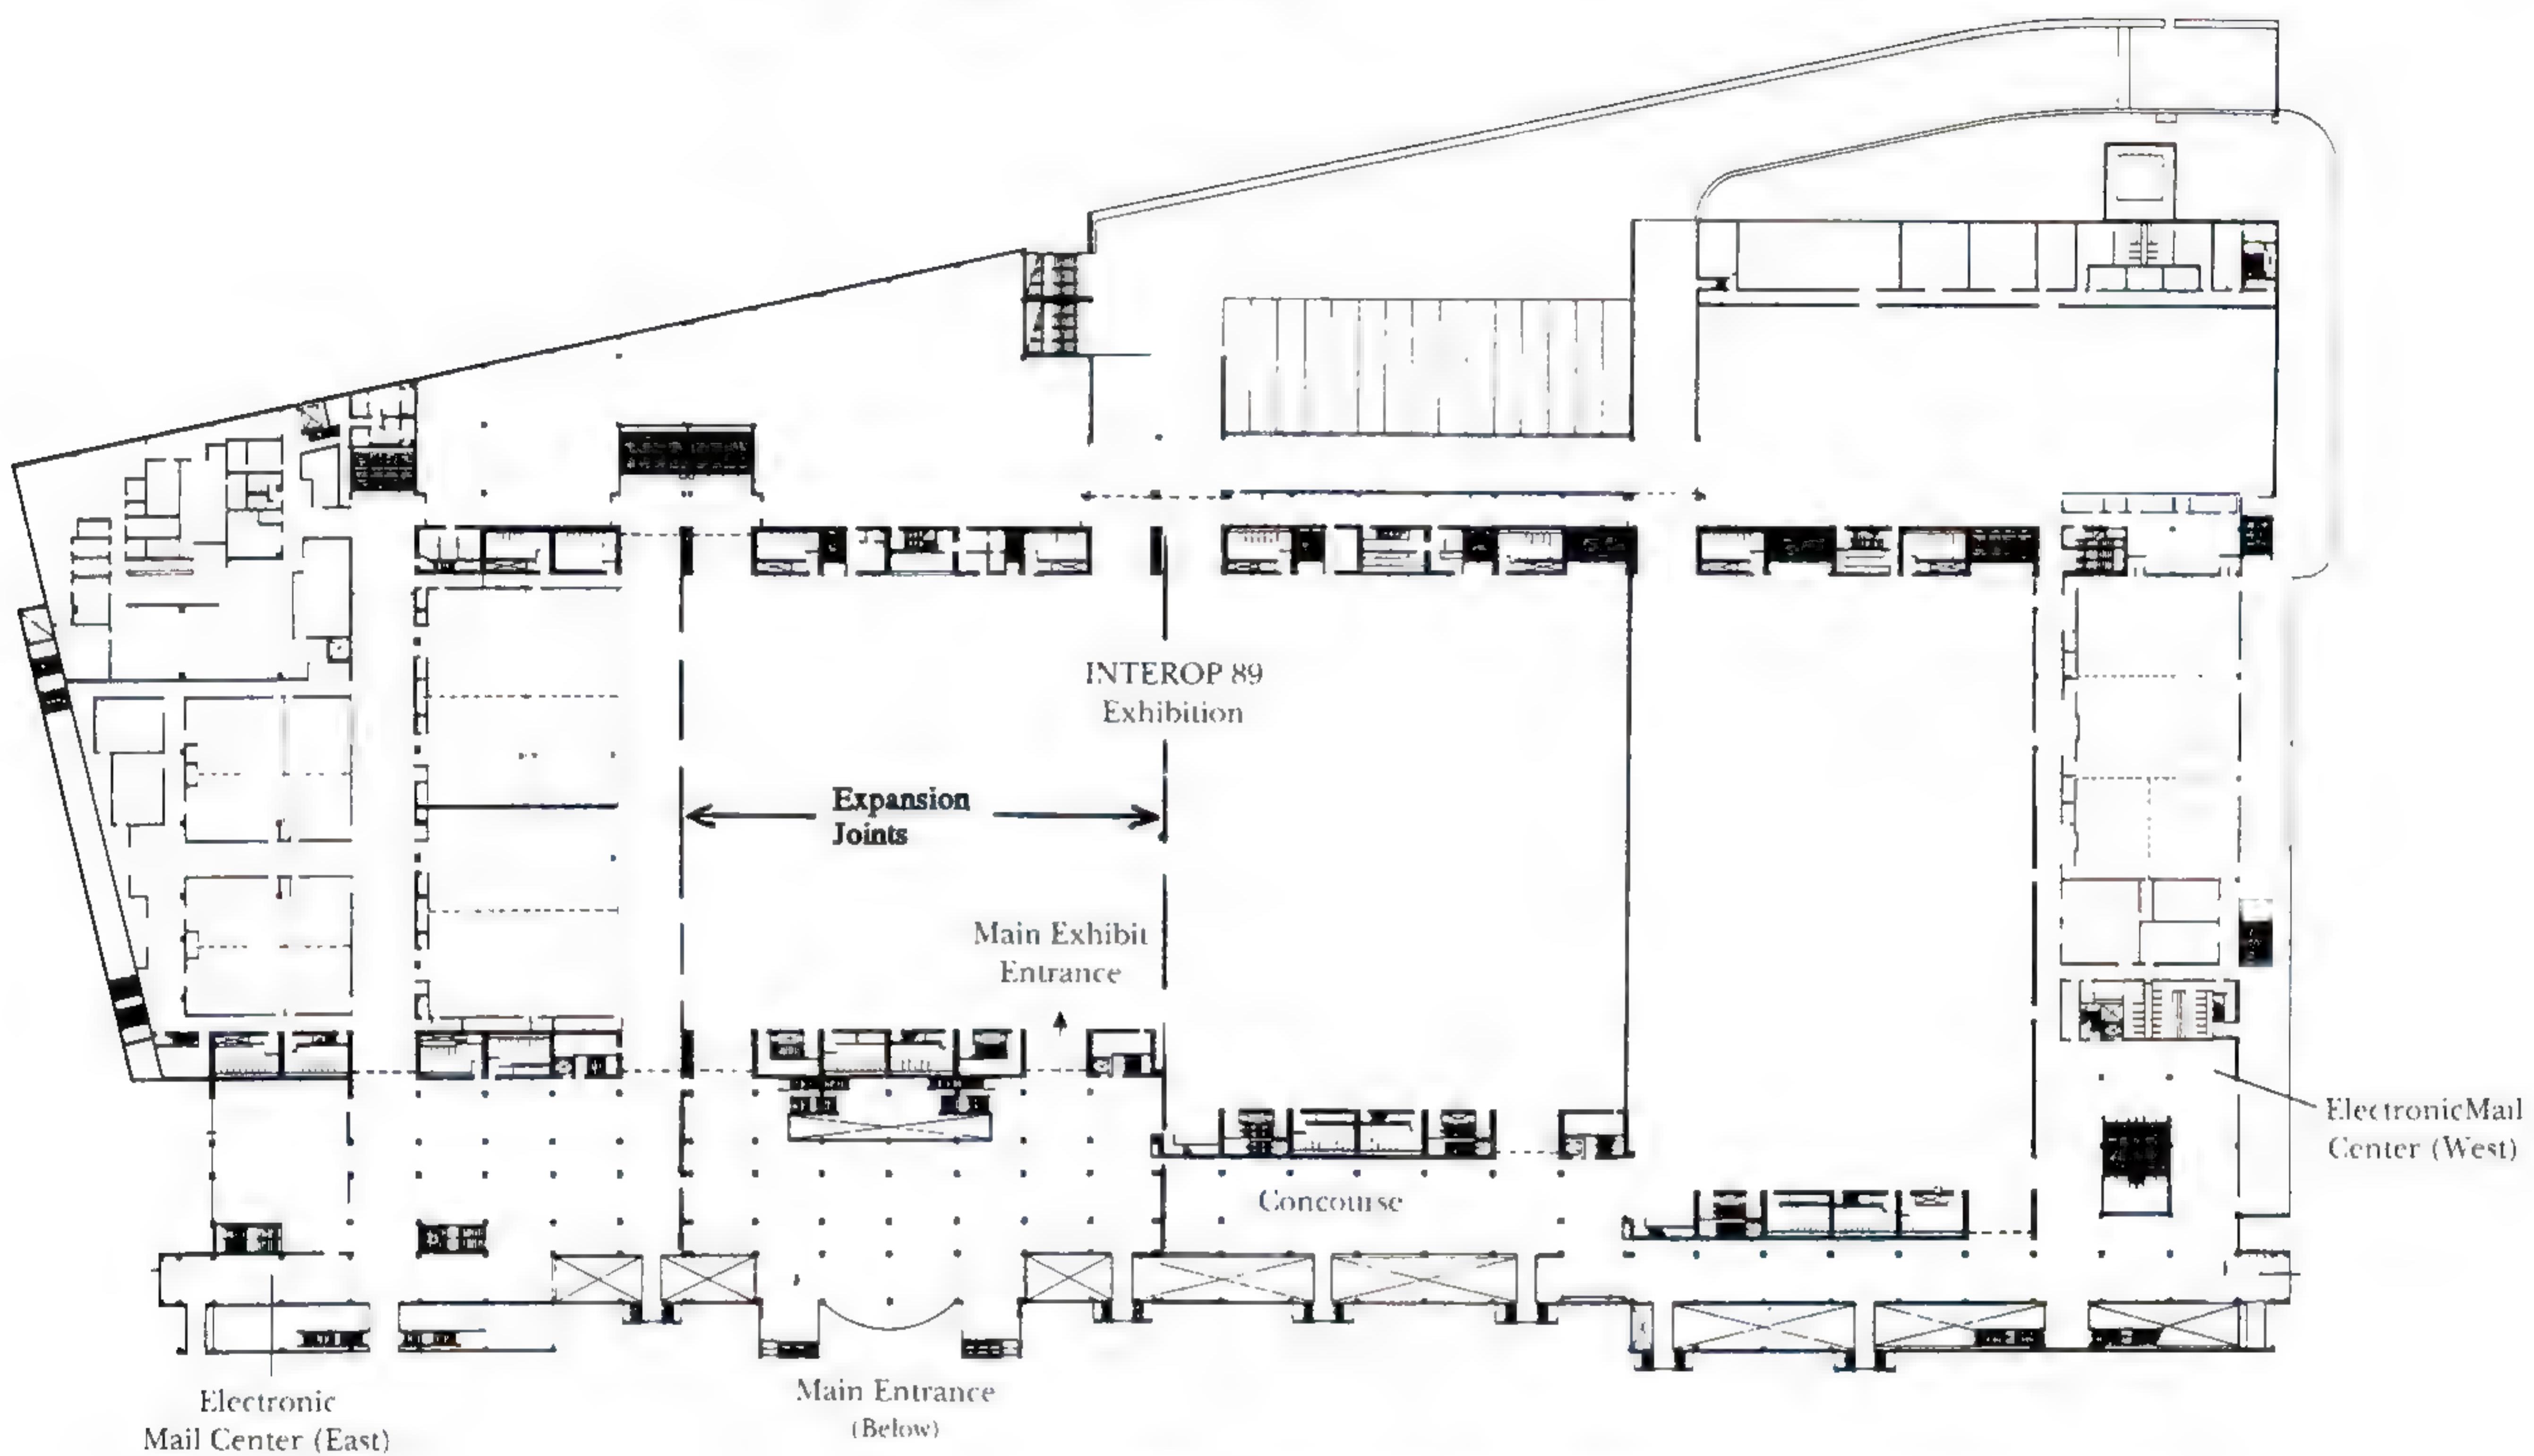


Figure 1: The San Jose Convention Center

Meanwhile, at the Fairmont hotel

At this point, we had a problem with the microwave connection to the hotel. The problem was getting from the room the receiving dish was in to the room the mail center was in. The hotel had told us they would be able to provide a connection between these rooms. We assumed they meant they would be able to connect us through the phone drops in the two rooms by jumpering them together in a phone closet. They thought all we needed to do was run modems between the two rooms, and they had refitted the rooms to have modular connectors in them. We decided that we had benefited from investing the time in sending someone to detail the convention center, so we sent someone over to look at the hotel. Their suggestion we went along with. It involved running a UTP line from the microwave room, to the stairwell, down to the floor the terminal room was on, and over. Fortunately, both rooms were on the same side of the hotel. This involved us using additional UTP transceivers, which we had over-ordered, but required an additional UTP fan-out unit, since you cannot connect UTP transceivers "back to back." Discussions with the SynOptics engineers resulted in them removing a unit from one of their demo racks and handing it over to us to use elsewhere.

continued on next page

The INTEROP 89 Network (*continued*)

The final problem with the mail centers, that resulted in us opening them late, after the INTEROP tutorials had started, was a shipping delay in the fiber we expected to run along the front of the buildings. We decided to wait for the fiber to show up because our shipping contacts told us it was in transit, and because we already had the appropriate hardware on hand to allow us to use it. After Monday, we decided that we were going to assume the fiber was never going to show up, and huddled together for another solution. The final decision was based solely on what we had left over. We ended up running thin Ethernet from the drop to both ends of the building for the mail centers. This introduced yet another problem. We had no equipment on hand to bridge from the thick Ethernet to the two thin Ethernets we were running. A quick talk to the engineers at the Wellfleet booth produced a box we could use as a bridge for the duration of the show. This was taken from the demo machines they had brought for use in their own booth.

Connectors

One serious problem we had was with the UTP connectors. These were standard eight conductor telco modular phone type connectors. Initially, the shipment of connectors we got could not be located. We sent out a runner to get more. One thing that people putting on something like this forget is the value of someone to go out and get the things you forget. The more you can keep your crews doing the stuff they do best, the more progress you make. Having someone who is willing to spend the day running around and getting parts, and lunch, for the crews is *very* important, and may be one of the most generally useful people you will have around. The runner returned with crimp-on connectors. Since I have seen these connectors before, I gave the crew starting on it a quick class in crimping the connectors with a pair of channel-lock pliers. They went off to be productive.

Figure 2: Part of the network installation team.

Left-to-right: Steve Knowles, Steve Larbig, Dave Bridgman, Glenn Trewitt, Phil Almquist, Peter de Vries, John Romkey, and Shelly de Vries.



Two types

Later they returned proclaiming a startling lack of success. As more people got involved, we discovered that there are *two* types of crimp-on connectors; one for stranded wire, one for solid. We, of course, had stranded connectors, and solid wire. The runner reported that the store he had bought the connectors at only had that one type. One of the people who had *extensive* experience showed us a connector he had in his pocket. It was much easier to use, since it required no special tools. It had a flared back, where one could lay the wire strands in channels, and screwing the cap on caused teeth to bite into the wires.

This proved to be much quicker than the crimp-on connectors (even after we had gotten some crimping tools), and much more reliable, but had the disadvantage that the part that stuck out of the wall was so wide that the fan-out boxes we used could not have connections in every jack, we had to leave empty jacks between them. While Peter had planned extra jacks, he had not planned on twice as many jacks. Fortunately, Synoptics was able to provide us with additional fan-out boxes for the duration of the show.

Here we got bitten by several problems. The connectors did not arrive on time, and the replacements that were bought were the incorrect type (even the packaging did not imply that there were two types). Fortunately, someone was able to put us in touch with a vendor who would extend us immediate credit for the connectors we needed and let us go and pick them up from them with no notice. Some of the crimping tools we had did not work correctly, and there were not enough of them if they had. The module we were using to test the lines was defective, since the adapter we plugged the cable into had a broken wire. Some of these things we could have taken care of before hand, like more crimping tools, and testing the tester out on a known good wire. Some of it, like the two types of crimp-on connectors, we should have known about, but most of us had never run into these problems before.

UTP or “All I got was a phone connection!”

One of the big problems was caused by people not paying enough attention. We ran *a lot* of UTP. UTP was used for several reasons:

- It was cheap, and easily available.
- It was easy to work with, and (we thought!) easy to connectorize.
- It didn't weigh much, so we could connect more runs to one anchor point on the ceiling.
- Using off-the-shelf technology, we could provide *either* thick or thin Ethernet access from it.

The topology was simple, and involved the use of several fan-out boxes. Fan-out boxes can be *cascaded* to a certain degree (cascading is hanging one fan-out box under another, instead of attaching them all at the same level). The destination end can either provide the user with a “thick” connection (really, a place to plug in a transceiver cable), or a “thin” connection (a place to attach a T-connector). Unfortunately, fewer people than we would have thought had heard of this technology. This would not be a problem, but people were being told that we had finished running wire around the show floor (they were concerned since their “roadie cases” were being moved in, filling the aisles, making it impossible for us to move cherry-pickers around the booths. People were constantly coming up to us telling us that they had a “phone” drop, but not the thick Ethernet drop they expected.

Very quickly we started putting up signs at the entrance doors telling people that the UTP drops above their booths were indeed network drops, and not to worry about it. We surrounded the INTEROP booth with these signs also. This did not seem to stem the tide very much. When people would come to ask a question, we would just point to the closest sign, and usually they would go away. Sometimes they would ask anyway. We made the mistake of telling them they would have “thick” Ethernet or “thin” Ethernet on the forms they filled out.

continued on next page

The INTEROP 89 Network (continued)

What we *should* have done was ask if they wanted a drop for a transceiver cable or a T-connector. Since we were vague about this, some people did not bring enough of their own equipment. Fortunately, a lot of companies brought more than they thought they could possibly use, and were very helpful in loaning connectors and cables to each other. Seeing competitors helping each other fix problems is unusual in todays competitive market. It is a shame that the customers that came to see the shonnet never got to see this wonderful cooperation.

Routing

Actually, there were few machines on the INTEROP show floor that misbehaved on the community wires. Some booths had routers between their networks and the shonnet internetwork. When possible, we got the network addresses they were using internally, and allowed them to be RIPed around the shonnet floor. Since the shonnet was connected to the outside world by several routes (including NASA Ames and the NSFNET NOC in Ann Arbor), we were very sensitive to routing issues. We allowed only one of the shonnet routers to peer with routers outside of our internetwork, and we only allowed it to advertise a route to the shonnet class B network. We were not interested in the dynamic routing on the NSFNET backbone finding out that there was a low cost route from the NSFNET NOC to NASA Ames across the shonnet.

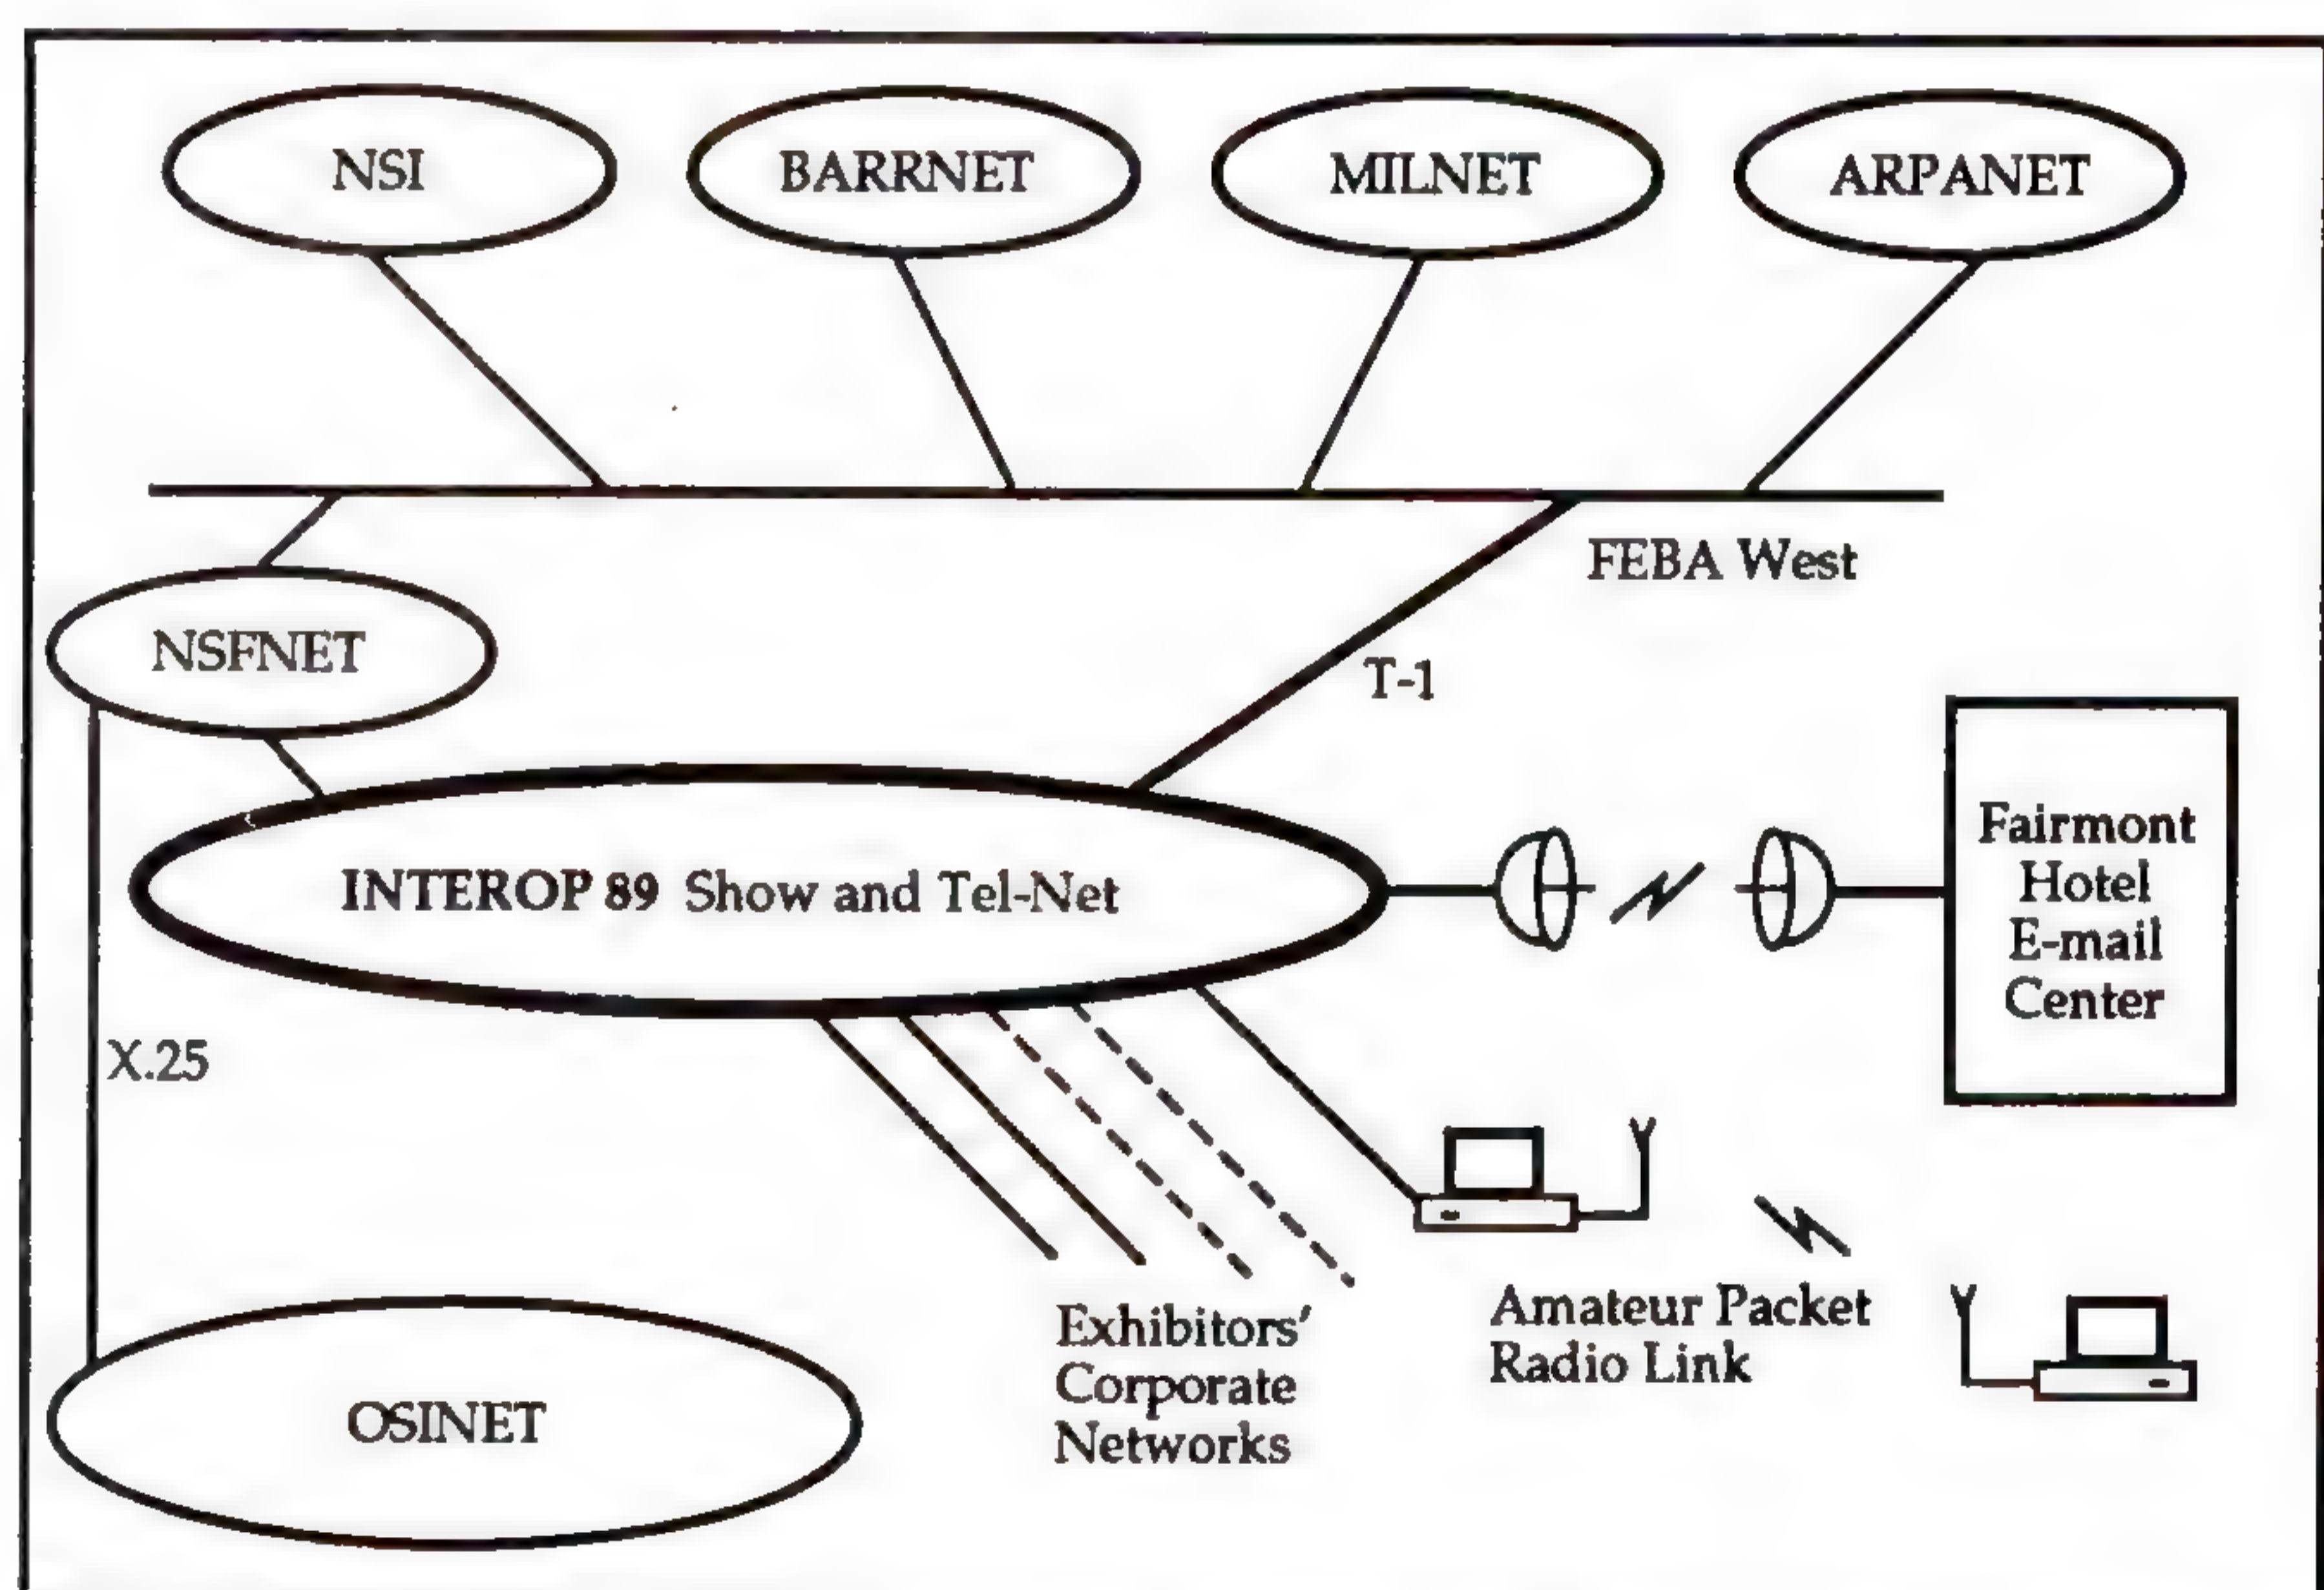


Figure 3: INTEROP 89 logical map

While this in itself would have been an impressive advertisement for the adaptability of TCP/IP to expand and contract its routes based on availability, we decided that it would be a hard thing to convey to the people visiting INTEROP 89, and it would also add to the load of what we thought would be an over-crowded backbone.

As it turns out, the shonnet backbone *was* very crowded. The network monitoring programs we had available were scrolling packets by faster than one could read. Fortunately, any of the routers we put on the shonnet backbone could be configured to broadcast routing information only on certain subnets, and accept information only from certain hosts.

This way, we did not have to worry about someone's mis-configured BSD-derived system trying to peer with the shonet backbone routers. Since there were several "fan-out" boxes in use, short term network outages were restricted to small physical parts of the much larger logical network, when one thin drop from a fan-out box went unterminated, usually because someone was adding another segment and T connector to it, the other drops in the fan-out box did not notice the outage, and continued to function as before.

Lessons learned

What makes an endeavor like this possible? The cooperation among several key people. People who know what you are trying to do very well, and are not so rigid that they cannot accept suggestions from out of their scope of knowledge. Maximizing communications is important. ACE provided the wiring crew with 8 walkie-talkies. These went far in speeding up the decision making process. We spent considerably less time tracking people down to collect their opinion about crises and planning changes than we would have had to otherwise. It allowed us to propagate changes to the crews faster, and to quickly reassign people to fight crises that were always springing up. Since we had several people around who, while not being knowledgeable about what we were trying to do, were willing to learn how to do the jobs we needed done, willing to follow orders, willing to go run out and pick up supplies, or call people to track down equipment, we freed other people to do more complex jobs.

Keeping people working on what they do best greatly increases the amount of work a crew can accomplish. Since the people assembled here all knew each other, we could quickly put together a good group to evaluate change proposals, or make suggestions about how to fix the problems we encountered. Often we found that even when incomplete groups met to make decisions, the groups would make a decision that the missing people would agree with. Some people find it hard to run a large project by consensus. We found it to be a process that speeded us up rather than slowed us down. Having several people with a complete "world view," who could be counted on to give an authoritative answer on 95% of the questions asked, allowed people trying to get work done to move faster since they could get their approval from a closer source.

Conclusion

All in all, this was a pleasurable experience. I got to work closely with some of the best people in our business, and I was able to learn a lot from the people around me. We ended up with something that the trade show attendees were excited about, and impressed with. I found the vendors in our industry to be made up of smart, helpful people willing to give their competitors a hand when they forgot equipment or cables needed to connect equipment to the shonet. While it was a very draining experience, I would be glad to do it again. We have already started thinking about INTEROP 90. From what I hear, since people thought it went so smoothly this time, we are going to have to make it more complex and flashy next year to impress the attendees. I can't wait to start playing with it!

STEV KNOWLES is a hacker at FTP Software. (This referring to the "Old School" of hackers, when being a hacker was a title bestowed upon you by others. It is not the only title bestowed on Mr. Knowles by others, but it is the only one we could get past the publisher.) Mr. Knowles had many diverse words (some of encouragement and some not) for the INTEROP 88 shonet engineering team, and therefore was drafted for the 89 team. He is currently involved in the IETF Benchmarking Methodology Working Group, and watches over several projects at FTP. He has no degrees, certificates, or warrants outstanding for his arrest. He can be reached by Internet e-mail as stev@ftp.com.



Introduction

A new feature at INTEROP 89, held in San Jose from October 2 through 6, 1989, was a demonstration of interoperability between OSI systems. Some thirteen vendors participated, demonstrating interoperation of such application protocols as X.400 and FTAM. Less visible, but just as important, was interoperation at the lower layers of the OSI stack, including the use of multi-protocol (IP and OSI) routers.

This article examines the INTEROP 89 OSI demo from the standpoint of the lower layers, particularly network layer routing and addressing. The architecture of the show network, the problems encountered, and the lessons learned are discussed.

Topology

The topology of the network was somewhat complex (Figure 1). The INTEROP Show Backbone network consisted of an Ethernet spine with a number of routers on it. Each router fed a number of Ethernets onto which hosts were attached. An OSI demo LAN was set up in the OSI booth, with connections to the show backbone, to OSINET via Accunet X.25, and to the rest of the Internet via NSFNET. OSI End Systems were located both on the demo LAN within the OSI booth and scattered around the show floor. Nearly all of the show backbone routers were switching OSI *Connectionless Network Layer Protocol* (CLNP) packets as well as DoD IP. Operation of CLNP over baseband (802.3), Token Bus (802.4), and Token Ring (802.5) LANs, as well as point-to-point links and X.25, was demonstrated. Connectionless network services were used exclusively.

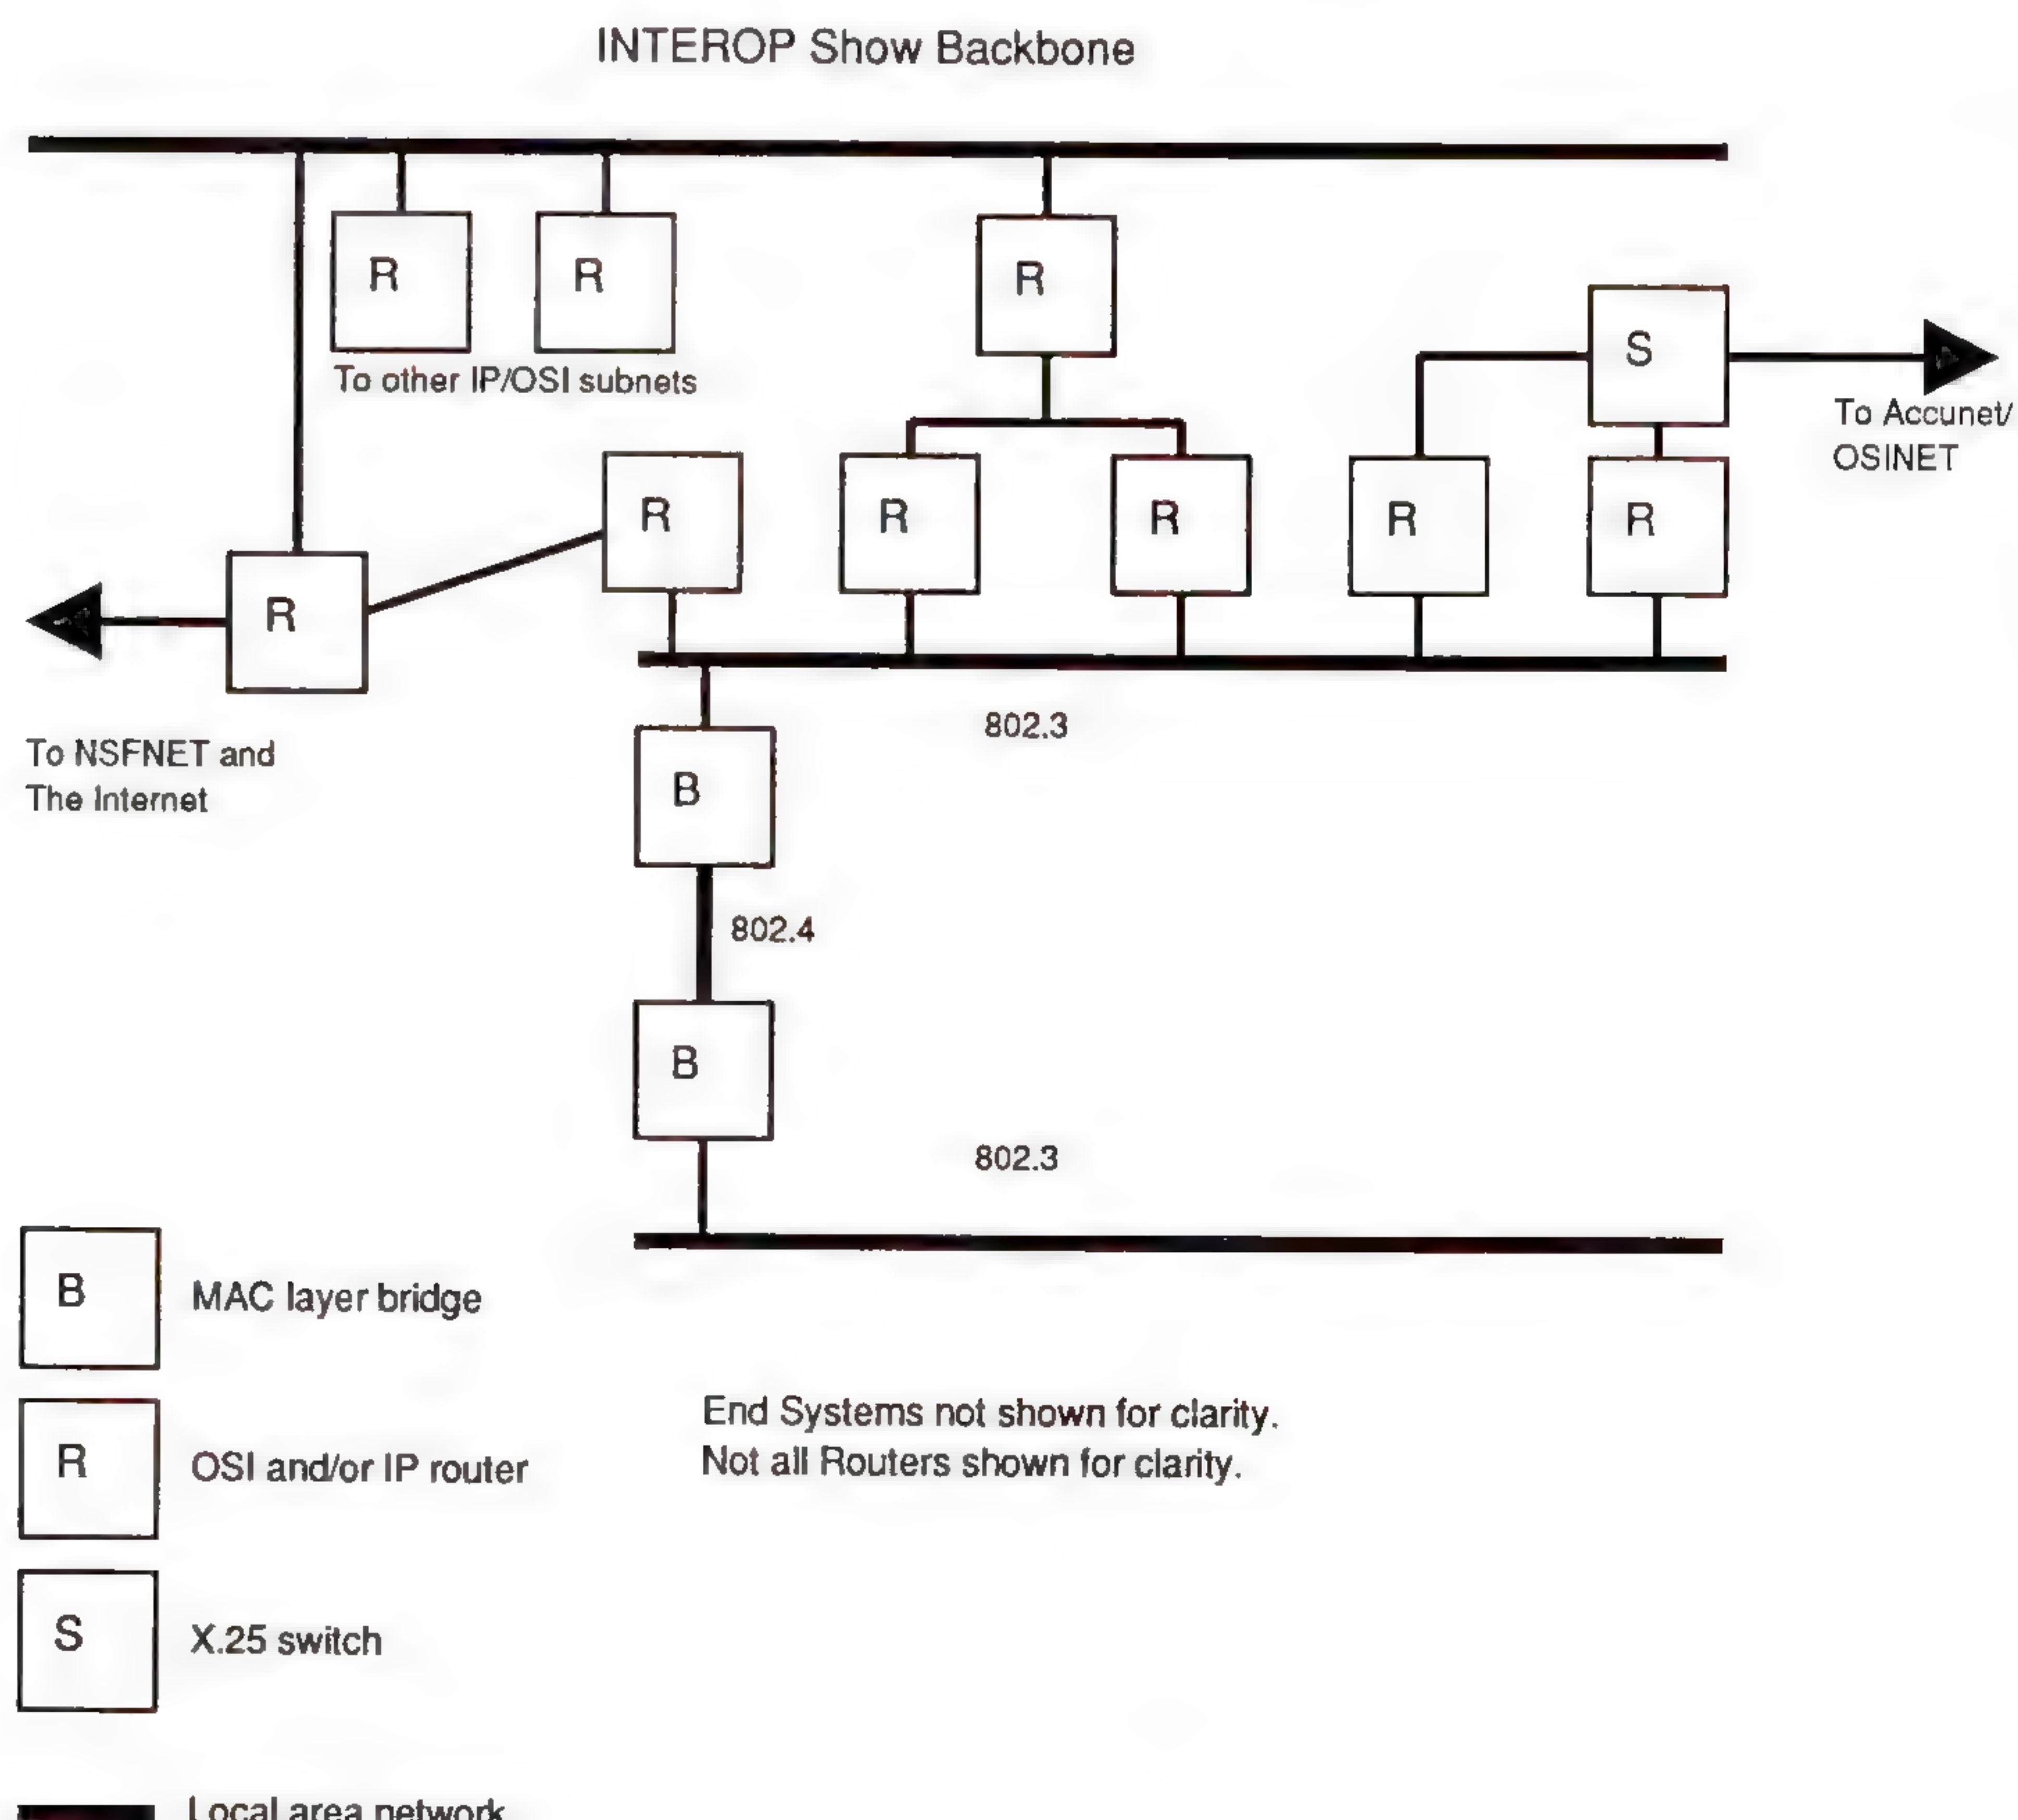


Figure 1: INTEROP 89 OSI Demo Topology

The OSI demo LAN consisted of two 802.3 LANs connected together with a pair of 802.3/802.4 bridges and an 802.4 broadband network. To this bridged LAN were connected five *Intermediate Systems* (ISs, routers) and a number of *End Systems* (ESs).

An OSI LAN in Ann Arbor, Michigan, was connected to the show network via NSFNET at a data rate of 1.536 Mbps. The path through NSFNET included 802.3, Token Ring, and point-to-point links.

A link to the Accunet public data network provided X.25 connectivity from the demo network to OSINET sites.

Addressing

OSI NSAP addresses were assigned in a hierarchical manner. Addresses had the following format:

47 0004 0046 RR SS NNNNNNNNNNNN SS

(a)	(b)	(c)	(d)	(e)	(f)	(g)
-----	-----	-----	-----	-----	-----	-----

- (a) Authority and Format Identifier (AFI). 47 means "ISO International Code Designator with Binary Abstract Syntax."
- (b) ICD 0004 = OSINET
- (c) OSINET Organizational ID 0046 = INTEROP Show Network
- (d)
and
- (e) OSINET Suborganization ID. This was assigned such that the first octet signified a router number within the show backbone, and the second octet selected a particular subnetwork (LAN).
- (f) SNPA (MAC) address
- (g) NSAP Selector

The choice of NSAP address format (from those outlined in ISO 8348 Ad. 2) was imposed by limitations in some End Systems (see below).

Routing

ES to IS routing was done, with one exception, using the ES-IS routing protocol. This protocol provides functionality similar to DoD ARP plus ICMP Redirect, and allowed End Systems and Intermediate Systems to discover one another dynamically.

IS to IS routing was set up using static tables. The hierarchical nature of the addresses made setting up these tables fairly easy. For instance, NSFNET nodes were set up such that all packets addressed to destinations starting with 47 0004 0046 05 were sent to Ann Arbor, and all other packets were sent to San Jose. Similarly, the show floor routers were configured to pass traffic to each other based on the Router ID field in the NSAP address, and packets with all other address formats were passed to the OSI demo LAN.

continued on next page

OSI at INTEROP 89 (*continued*)

Problems	<p>Not surprisingly, a number of problems cropped up during the course of the show, ranging from the mundane to the bizarre.</p> <p><i>LANs Bridged Together:</i> The OSI demo LAN was bridged at the MAC layer to one of the show subnetworks without the knowledge of the network architects. This led to a number of problems, including routing loops. It is interesting to note that OSI networks fare far better in this situation than IP networks. In fact, if the ANSI IS-IS routing protocol were used and the LANs bridged together were members of the same routing area, bridging would cause no problems at all. Difficulties in the DoD IP world with renegade bridges are due to the presence of Proxy ARP and to the fact that IP addresses have explicit subnet addresses embedded within them.</p> <p><i>Bad Static SNPA <-> NSAP Mappings:</i> One End System vendor did not support ES-IS. As a result, static NSAP <-> SNPA mappings had to be entered into all other ESs and ISs. Since each table entry consisted of some forty hexadecimal digits, the probability of mistyping them was quite high. Several systems were misconfigured, temporarily resulting in poor connectivity for the vendor in question.</p> <p><i>Bad Static Routes:</i> Because routes, albeit simple ones, had to be entered manually, and because not all of the router vendors conferred before the show, problems occurred with some of the routes. Not only did machines on different subnetworks have difficulty communicating, but so also did machines sitting next to each other on the same LAN. Machines on the demo LAN suffered from Probabilistic Black Holes, causing communications between machines to work some of the time and fail at other times.</p> <p>Probabilistic Black Holes happen because of the way local routing works on OSI LANs. When an End System wishes to send a packet for which the destination's SNPA address is unknown, it instead forwards it to an Intermediate System and expects to receive a redirect if that IS is not the right choice (such as when the destination is on the same LAN). The choice of IS is a local decision; some implementations use a round robin scheme, some use the last IS from which an ES-IS protocol "IS Hello" was received, etc. If one or more of the ISs are not routing properly, ESs unlucky enough to pick the wrong IS will be unable to communicate with other systems.</p> <p>Routes were eventually straightened out by carefully examining each router and working with the vendors to correct inconsistent configurations.</p>
Bugs	<p>Due to the prototype nature of some of the routers, some bugs were discovered. Most of these were trivial; a couple will require some new code. In general the routers worked well.</p>
Lessons learned	<p>A number of useful lessons learned in the process of setting up a real multivendor OSI production environment are worth sharing.</p> <p><i>Use the ES-IS Routing Protocol:</i> This should go without saying, but one vendor did not support it at the time of the show.</p>

This caused direct problems, such as the misconfiguration of static mapping entries, and indirect problems, such as the requirement that SNPA addresses be embedded in NSAP addresses, which made address administration more difficult.

Use an IS-IS Routing Protocol: Even in the relatively simple routing environment of INTEROP, black holes and routing loops were experienced. Most routing problems could have been avoided with the use of a dynamic routing protocol. This will be a fact of life until vendors feel that the ANSI IS-IS protocol is stable enough in its progression toward standardization to be implemented.

Don't Listen To Your Own IS Hellos! One implementation had hardware that received its own broadcasts and multicasts. This caused routing loops when it decided to forward packets to the last IS it heard a Hello from (itself).

ESs Should Allow Arbitrary NSAP Addresses: There is no reason that an End System should know or care about the structure of its own (or any other) NSAP address, outside of the position of the Selector field (the last octet). Addressing at INTEROP was constrained because of the assumptions made about addresses by some End System vendors. End Systems should treat NSAP addresses as opaque octet strings and should be configurable to any address.

OSI Subnetworks are not IP Subnets: In the DoD IP world, the word “subnet” has two related meanings. It refers to that piece of communications infrastructure that provides a “single hop” from the standpoint of network addressing (such as a LAN or the ARPANET). It also refers to a portion of the IP address itself. IP addresses contain the address of the wire to which the system is connected (which is why IP routers need a different address on each interface).

In the OSI world, the word “subnetwork” refers only to the medium—the NSAP address is not inherently related to the wire. NSAP addresses are hierarchical, but the hierarchy ends at a point higher than the subnetwork. In the address format called out by the ANSI IS-IS routing proposal, the NSAP address is hierarchical down to the address of a particular Area (which may consist of multiple subnetworks), but no further. This has a number of implications:

- Only one address (Network Entity Title) per ES or IS
- All ISs in an area must keep a complete list of ESs and ISs in the area
- Routing is not tied to low level topology (no problems with extra bridges)
- Address administration is easier (and can be automatic)
- No “local/distant” routing decision by ESs is necessary (or possible—no subnet masks, default routers, IGP eavesdropping, etc.)
- ISs can deterministically report “Host Unreachable”

continued on next page

OSI at INTEROP 89 (*continued*)

Although NSAP addresses at INTEROP were structured to contain subnetwork addressing, and End Systems had their addresses assigned that way, Intermediate Systems could not be so addressed, due to the NET advertised by an IS serving a particular LAN often did not have the same "subnet" number as the End Systems on that LAN. At least one ES was confused by this.

At least two routers did not handle all of the implications of all this properly. In particular, an issue exists with packets addressed to unreachable End Systems. Some IS must be able to make the final decision about whether an ES is reachable or not. When the system was reachable, the static routing information eventually caused the packet to arrive at an IS that had heard an ES Hello for the destination. However, when the destination was unreachable, the final IS had not heard of the ES, and, because it did not have any sort of entry that said "packets with this NSAP prefix are destined for this subnetwork," forwarded the packet to a default router, which, having a entry pointing back at the final IS, forwarded the packet back, and so forth.

In an environment running the ANSI IS-IS protocol, all routers within the destination area have a complete list of ESs in the area, and thus can determine quickly whether the destination is reachable or not. In the absence of an IS-IS protocol, care must be taken to ensure that an IS knows the NSAPs that it is supposed to be responsible for so that it can either forward or reject packets as necessary. Note that this requires a tighter binding between address structure and topology than would otherwise be necessary, but is a fact of life so long as ISs don't have a mechanism to redistribute the addresses of attached ESs.

Echo is useful: Some of the routers provided implementations of an equivalent to the ICMP Echo ("Ping") function. Such functionality is not in the CLNP standard, and ad hoc solutions were provided. The Echo function proved quite useful in tracking down routing difficulties and the like. It would be most useful if such a function were to become part of the standard!

Improve planning for lower layers: Many of the INTEROP OSI demo participants did not appear to be particularly interested in routing and addressing, and this ultimately contributed to the difficulties experienced. Although a significant and impressive effort was made by the vendors to plan for the show, set up a staging area, and so forth, most of the technical discussions dealt with the applications and not with the infrastructure. The Network Layer is only invisible when it works!

Conclusion

For the most part, the problems experienced during the demonstration were caused by misconfiguration or minor bugs. Very few fundamental flaws were seen in either the ESs or the ISs.

Although a small scale demonstration such as this certainly does not prove the soundness of the protocols and implementations in a large network, it is nonetheless heartening to see a significant number of different vendors bring their equipment together and have it work well.

On the whole, the OSI Interoperability demo at INTEROP 89 was a success, and the experience gained in this endeavor should help set the stage for the emergence of an OSI Internet in the near future.

Glossary

802.3 A local area network (LAN) technology using Carrier Sense Multiple Access with Collision Detection (CSMA/CD) as the access method. IEEE 802.3 is essentially equivalent to Ethernet, although it uses slightly different framing. 802.3 and Ethernet traffic can run on the same wire.

802.4 A LAN technology using token passing as the access method on a bus topology ("Token Bus").

802.5 A LAN technology using token passing as the access method on a ring topology ("Token Ring").

ANSI The American National Standards Institute. An organization that accredits standards committees in the United States. Such accredited standards committees create American National Standards and may be affiliated with related international committees in ISO.

Black Hole A routing phenomenon in which transmitted packets disappear without a trace. Typically caused by routing inconsistencies.

CLNP The OSI Connectionless Network Protocol, ISO 8473. Roughly the OSI equivalent of DoD IP. [See *ConneXions*, Vol. 3, No. 10].

ES End System. A system on which applications run. All layers of the OSI stack are implemented on End Systems. Roughly equivalent to an IP Host.

ES-IS The End System–Intermediate System routing protocol, ISO 9542. This protocol is used on LAN media to allow End Systems and Intermediate Systems to discover one another and to map network (NSAP) addresses to media (SNPA) addresses. It also provides the means for Intermediate Systems to provide End Systems with optimal routes. The ES-IS protocol provides the same functionality, more or less, as the ARP and ICMP Redirect protocols in the DoD IP protocol suite. [See *ConneXions*, Vol. 3, No. 8].

FTAM File Transfer, Access, and Management. An OSI application protocol providing file transfer and manipulation services.

IS Intermediate System. A system which provides the services necessary to allow End Systems to exchange information. Intermediate Systems implement only the lower layers (Physical, Data Link, Network) of the OSI stack. Roughly equivalent to an IP router.

IS-IS An Intermediate System–Intermediate System routing protocol. Such a protocol allows Intermediate Systems to exchange routing information in order to dynamically establish routes between systems. One such protocol (the "ANSI IS-IS Protocol") is being progressed through ISO as a proposed standard. In IP parlance, such protocols are known as IGPs (Interior Gateway Protocols). [See *ConneXions*, Vol. 3, No. 8].

continued on next page

OSI at INTEROP 89 (*continued*)

ISO International Organization for Standardization. An international standards body made up of representatives of the national standards bodies of the member nations. [See *ConneXions*, Vol. 3, No. 12].

NET Network Entity Title. The network address of a system, particularly an Intermediate System. Basically equivalent to an NSAP address, except that Intermediate Systems have no transport layer, so they do not have a network/transport layer boundary on which an NSAP could reside. NETs do not address any higher layer service. Basically equivalent to an IP address.

NSAP Network Service Access Point. A conceptual point on the network/transport layer boundary in an End System at which transport services are provided. An End System may have multiple NSAPs. The address of an NSAP (the NSAP Address, or NSAPA) is globally unique and is used by the network layer for addressing and routing. The final octet of the NSAP address is known as the NSAP Selector, which selects one of a number of NSAPs within the system. Equivalent to the combination of an IP address and the IP protocol type (not to be confused with the TCP/UDP port number).

OSI Open Systems Interconnection. An architecture describing the structure of Open Systems, which are capable of communicating in a heterogeneous, multivendor environment. The term "OSI" is also used to describe the suite of protocols that provide the services defined by the OSI model.

SNPA Subnetwork Point of Attachment. The conceptual attachment point of a system on a subnetwork. The SNPA address is used to access a particular system on a multipoint subnetwork. On a LAN, the MAC address is used as the SNPA address.

Subnetwork A communications medium that provides a "single hop" from the standpoint of the network layer. This can be a point-to-point link, a LAN, or an X.25 network, for example.

X.400 The OSI Message Handling Service. A protocol that provides document and electronic mail exchange. [See *ConneXions*, Vol 3, No. 5].

DAVE KATZ is an Internet Engineer on the NSFNET Backbone project for the Merit Computer Network. He is active in standards activities in both the DoD Internet (the Internet Engineering Task Force) and OSI (ANSI X3S3.3, OSI Network/Transport layers). Dave is an agnostic when it comes to network religion.

[Ed.: See also "INTEROP 88 Conference Report," in *ConneXions*, Volume 2, No. 11, November 1988, "The INTEROP 88 Network—behind the scenes," in *ConneXions*, Volume 3, No. 2, February 1989, and "Highlights from INTEROP 89," in *ConneXions*, Volume 3, No. 11, November 1989].

Mark your calendar: *INTEROP 90 will be held in San Jose, California, October 8–12, 1990.*

Letters to the Editor

Ole,

In a recent *ConneXions* article entitled "LAN Packet Demultiplexing" (Volume 3, no. 12, December 1989), James VanBokkelen stated: "Regrettably, neither [Xerox nor the IEEE] has ever made the list of assigned [Ethernet] types public...".

Xerox Numbers Administration

This statement is misleading, at best. Earlier this year, I was in need of such information. A call to the Network Systems Administration Office of the Xerox Systems Institute yielded just such a list. Those requiring a copy may find it helpful to request the "Public Ethernet Packet Types" memo (19 May 1988) from their "Numbers Administration" organization. Their phone number is 408-737-7900.

Some companies have requested that their assigned private numbers not be published. For information on these types, each company must be contacted directly. This is clearly not a shortcoming on the part of Xerox. I found them to be most helpful.

*Lou Steinberg
NSFNET Software Development
IBM University and College Systems*

Dear Ole,

Length field

I read with great interest the article entitled "LAN Packet Demultiplexing" in the December 1989 issue of *ConneXions*. It contains a serious error which should be corrected before it spreads. The section entitled "IEEE Packet Header" neglects pointing out that the two bytes used for the Ethernet packet type are interpreted as a length field in an 802.2 header. With one small exception (PUP), the legal values for lengths and Ethertypes are disjoint, and hence typically can and do coexist, and be demultiplexed on one physical cable. Reading this article as written leaves one with the impression that the two bytes used for Ethertype are used for the DSAP and SSAP values. Saying that the 802.2 header is "normally either 3 or 4 bytes long" contributes to this confusion, since it doesn't include the 2 bytes length field.

Yours truly,
*Harry J. Saal, President
Network General Corp.*

The Author responds:

Regarding Mr. Steinberg's point: Perhaps I should have said "...complete list of assigned types..." Comparison of the Xerox list with those maintained by the network community will give an estimate of the relative numbers of "public" and "private" types. In any case, a lot of water has passed under the bridge since May 1988.

Regarding Dr. Saal's comments: The IEEE 802.2 document defines the "LLC protocol data unit" as starting with the Destination SAP. The two bytes that were redefined (from Ethertype to length) are part of the 802.3 MAC header. There is no corresponding field in the 802.5 MAC header, where the media doesn't impose a minimum packet length. Demultiplexing is only a small part of "Ethernet vs. 802.3"; the topic is broad enough that it deserves treatment in a separate article, covering both software and hardware. —James VanBokkelen

Book Review

ISDN Design: A Practical Approach by Steve Hardwick, Academic Press, 1989. ISBN 0-12-324970-8, 152 Pages, Hardbound.

Introduction

**ISDN will give us
cheap WANs**

**A book for the
ISDN product designer**

One of the most notable items about the forthcoming *Integrated Services Digital Network* (ISDN) is the dearth of informative publications for the designer. Much of the material available is either utterly devoid of any functional information, preferring to wax glowingly on the telecommunications paradise of what an end-to-end digital telephone network will mean (as if we did not know already), or, in the opposite case, is loaded with such arcane technical information and mnemonics to the point of being where you can't find the forest for the trees. (A first requirement for any text is to put knowledge into a structure). Even those of us who have labored to understand ISDN over the years would rather hike five miles than crack open another illiterate BELLCORE tome, so it is no wonder that the topic of ISDN is not widely understood.

The major impact of ISDN on the LAN world will be the transformation of the telephone network into a ubiquitous and cheap WAN (discount ARPANET!). Part of the reason why this has not yet been exploited is that developers and businesses (not to mention the RBOC's themselves) don't clearly understand the value ISDN offers. Another has been the incompetent attempt to market this technology as a replacement for LANs (who wants a 64Kb/s LAN?). Finally, there are the "bandwidth bigots," those who believe that any communications technology is not any good at less than 1Mb/s or higher, completely forgetting that the ARPANET backbone that many of us were weaned on was 56Kb/s technology. To this day, more than half of all long distance communications are done at bit rates less than 64Kb/s. In sum, poor valuation and marketing of this technology, coupled with a complete ignorance of the history and evolution of telecommunications and LAN technologies, has led many to misjudge or misunderstand where ISDN sits in the telecom and data-com world.

This book is gratefully in the category of first, a clearly written work, and second, on ISDN Product Design. The book is written expressly for the intended designers of ISDN-compatible equipment and systems, and is also a good first reference to be consulted after the decision to support ISDN has been made ("now what do I do?"). Those who are looking for a book that skirts the technical background might be a little taken aback by the discussions of eye patterns and jitter that are mentioned as early as chapter two, but the serious reader will be rewarded by the clear descriptions provided throughout of the telecommunications terms and concepts that are the heart of ISDN. As the book was written by a hardware engineer, it can't help but mention design hints for board layout, decoupling, and signal parameters, but the book is not oriented around these areas, as indicated by its OSI-oriented descriptions of ISDN protocol layers.

The author frequently stresses that ISDN is a multidisciplinary technology; to accomplish an understanding of ISDN one needs to learn a little telephony, computer networking, digital transmission, and engineering.

As you might expect, the topics covered in the book meander through signal diagrams on one page, then software functional diagrams on the next, and perhaps followed up by a historical note of how a part of the telephone system came to be a certain way (like the evolution of T-1 signalling). Even if you have no interest in ISDN, this book is valuable for its broad coverage of telephony/telecommunications nomenclature in a succinct style.

Unfortunately, the book's index is small and does not contain a complete set of references to all the key phrases covered. This would be a greater problem if the book was larger, but for now it is only a minor nuisance. Other problems include the lack of coverage of the long distance "U" interface, and specifics on North American 2B1Q and Primary Rate interfaces. Some of this is understandable due to the fact that the standards were settling when the book went to type. However, it is a significant omission as it renders the book out of date (even though it is the most up to date book available!).

Software designers will be disappointed by the lack of discussion on network layer software. Similarly, systems designers will find that the book runs out of steam in the "Case Study of a Terminal Design," which is given a mere fraction of a page of text. Both these sections are too light in the treatment of the subject. As a consequence, a good treatise on ISDN software design and integration has yet to be written.

Recommended

In summation, *ISDN Design* is a good practical guide to the nuts and bolts of designing a first ISDN product. ISDN is similar in many respects to network management; simple in theory but fiendish in detail. Steve Hardwick, an authority on ISDN for many years, has written a very readable book which does much to structure the process of understanding ISDN. Any engineer, faced with the need to build an ISDN interface, would value this book.

—Bill Jolitz

RFC on the Internet Worm now available

RFC 1135: The Helminthiasis of the Internet is now available from the Network Information Center in the online library at host NIC.DDN.MIL.

"The obscure we see eventually, the completely apparent takes longer." —Edward R. Murrow

This memo takes a look back at the helminthiasis (infestation with, or disease caused by parasitic worms) of the Internet that was unleashed the evening of 2 November 1988. This RFC provides information about an event that occurred in the life of the Internet. This memo does not specify any standard. Distribution of this memo is unlimited.

This document provides a glimpse at the infection, its festering, and cure. The impact of the worm on the Internet community, ethics statements, the role of the news media, crime in the computer world, and future prevention is discussed. A documentation review presents four publications that describe in detail this particular parasitic computer program. Reference and bibliography sections are also included.

—Joyce K. Reynolds

CONNEXIONS

480 San Antonio Road
 Suite 100
 Mountain View, CA 94040
 415-941-3399
 FAX: 415-949-1779

FIRST CLASS MAIL
 U.S. POSTAGE
 PAID
 SAN JOSE, CA
 PERMIT NO. 1

CONNEXIONS

PUBLISHER Daniel C. Lynch

EDITOR Ole J. Jacobsen

EDITORIAL ADVISORY BOARD Dr. Vinton G. Cerf, Vice President, National Research Initiatives.

Dr. David D. Clark, The Internet Architect, Massachusetts Institute of Technology.

Dr. David L. Mills, NSFnet Technical Advisor; Professor, University of Delaware.

Dr. Jonathan B. Postel, Assistant Internet Architect, Internet Activities Board; Division Director, University of Southern California Information Sciences Institute.

Subscribe to CONNEXIONS

U.S./Canada	\$125. for 12 issues/year	\$225. for 24 issues/two years	\$300. for 36 issues/three years
--------------------	---------------------------	--------------------------------	----------------------------------

International	\$ 50. additional per year (Please apply to all of the above.)
----------------------	--

Name _____ Title _____

Company _____

Address _____

City _____ State _____ Zip _____

Country _____ Telephone () _____

Check enclosed (in U.S. dollars made payable to CONNEXIONS).

Charge my Visa MasterCard Am Ex Card # _____ Exp. Date _____

Signature _____

Please return this application with payment to: CONNEXIONS

Back issues available upon request \$15./each
 Volume discounts available upon request

480 San Antonio Road Suite 100
 Mountain View, CA 94040
 415-941-3399 FAX: 415-949-1779

CONNEXIONS